



Advisory Editorial Board:

DOMINIK VAN AAKEN
FREDERIK AHLEMANN
CHRISTOPH BODE
ROLF BRÜHL
JOACHIM BÜSCHKEN
LEONHARD DOBUSCH
RALF ELSAS
DAVID FLORYSIK
GUNTHER FRIEDL
WOLFGANG GÜTTEL
CHRISTIAN HOFMANN
KATJA HUTTER
LUTZ JOHANNING
STEPHAN KAISER
ALFRED KIESER
NATALIA KLIEWER
DODO ZU KNYPHAUSEN-AUFSEß
SABINE T. KÖSZEGI
ARJAN KOZICA
TOBIAS KRETSCHMER
HANS-ULRICH KÜPPER
REINER LEIDL
ANTON MEYER
MICHAEL MEYER
GORDON MÜLLER-SEITZ
J. PETER MURMANN
BURKHARD PEDELL
MARCEL PROKOPCZUK
TANJA RABL
SASCHA RAITHEL
ASTRID REICHEL
KATJA ROST
MARKO SARSTEDT
DEBORAH SCHANZ
ANDREAS G. SCHERER
STEFAN SCHMID
UTE SCHMIEL
CHRISTIAN SCHMITZ
PHILIPP SCHRECK
GEORG SCHREYÖGG
LARS SCHWEIZER
DAVID SEIDL
THORSTEN SELLHORN
ANDREAS SUCHANEK
ORESTIS TERZIDIS
ANJA TUSCHKE
SABINE URNIK
STEPHAN WAGNER
BARBARA E. WEIßENBERGER
ISABELL M. WELPE
HANNES WINNER
CLAUDIA B. WÖHLE
THOMAS WRONA
THOMAS ZWICK

JUNIOR MANAGEMENT SCIENCE

Layla Martin , Extending Kolkata Paise Restaurant Problem to Dynamic Matching in Mobility Markets	1
Estevan Vilar , Word Embedding, Neural Networks and Text Classification: What is the State-of-the-Art?	35
Nicholas Herold , The Tax System and Corporate Payout Policies	63
Frédéric Herold , Personal Taxes and Corporate Investment	81
Christiane Czech , Living is Easy with Eyes Closed – Strategische Unwissenheit und eigennütziges Verhalten	101
Silie Homayon Nawabi , The Effect of ECB's Corporate Sector Purchase Programme on CDS Premia – An Empirical Analysis	123



Extending Kolkata Paise Restaurant Problem to Dynamic Matching in Mobility Markets

Layla Martin

Technische Universität München

Abstract

In mobility markets – especially vehicle for hire markets – drivers offer individual transportation by car to customers. Drivers individually decide where to go to pick up customers to increase their own utilization (probability of carrying a customer) and utility (profit). The utility drivers retrieve from customers comprises both costs of driving to another location and the revenue from carrying a customer and is thus not shared between different drivers. In this thesis, I present the Vehicle for Hire Problem (VFHP) as a generalization of the Kolkata Paise Restaurant Problem (KPRP) to evaluate different strategies for drivers in vehicle for hire markets. The KPRP is a multi-round game model presented by [Chakrabarti et al. \(2009\)](#) in which daily laborers constitute agents and restaurants constitute resources. All agents decide simultaneously, but independently where to eat. Every restaurant can cater only one agent and agents cannot divert to other resources if their first choice is overcrowded. The number of agents equals the number of resources. Also, there is a ranking of restaurants all agents agree upon, and no two resources yield the same utility. The VFHP relaxes assumptions on capacity and utility: Resources (customers) are grouped in districts, agents (drivers) can redirect to other resources in the same district. As the distance between agent and resource reduces the agent's utility and the location is not identical for all agents, the utility of a given resource is not identical for all agents. To study the impact of the different assumptions, I build four different model variants: Individual Preferences (IP) replaces the shared utility of the KPRP with uniformly distributed utilities per agent. The Mixed Preferences (MP) model variant uses the utility assumption of the VFHP, but the capacity of all districts remains 1. The Individual Preferences with Multiple Customers per District (IPMC) model variant groups customers in districts, and uses the uniform utilities introduced in the IP model variant. Mixed Preferences and Multiple Customers per District (MPMC) implements all assumptions of the VFHP. In this thesis, I study different strategies for the KPRP and all variants of the VFHP to build a foundation for an incentive scheme for dynamic matching in mobility markets. The strategies comprise history-dependent and utility-dependent strategies. In history-dependent strategies, agents incorporate their previous decisions and the utilization of resources in previous iterations in their decision. Agents adapting utility-dependent strategies choose the resource offering the highest utility with a given probability.

Keywords: vehicle for hire markets; distributed decision making; agent-based modelling; congestion game; limited rationality

1. Introduction

Mobility markets, or in particular vehicle for hire markets, comprise all modes of shared, but individual transportation with a driver, in particular with a short-term focus (e.g., taxis, Lyft, and Uber). In mobility markets, drivers individually decide where to look for customers. However, the average idle time of taxis is about 25–50%

in most cities where data is available ([Linne+Krause Marketing-Forschung, 2011](#); [Cramer and Krueger, 2016](#); [Linne+Krause Marketing-Forschung, 2016](#)). Though excess capacity can partially explain these numbers, utilization could be increased, if drivers would be distributed across the city more efficiently. In contrast to underutilization, passengers have to wait for more than 20 minutes in approximately every third case in other cities ([Rayle et al.,](#)

2014), suggesting that the drivers are not at the locations where they are needed.

To address these inefficiencies in vehicle for hire markets, coordinators could instruct drivers where to wait for customers. In current business models, however, this is not possible, since drivers are not employees of the coordinators. Hence, they try to maximize their individual profits by deciding independently where to look for customers without considering the social welfare or utilization of other agents. In practice, there are approaches like 'surge pricing' (price adapts dynamically to changes in demand and supply with the goal to influence demand and supply, e.g. increase supply by increased price) to respond to expected peaks in demand, though literature on the efficiency of different driver strategies is limited (Chen and Sheldon, 2015; Hall et al., 2015; Rogers, 2015). One, therefore, has to turn the attention to the coordination amongst drivers: Drivers maximize their individual utility, but their utility inversely depends on the number of agents selecting the same option. Thus, drivers benefit if there are less other drivers in the same district than available customers, thus, deciding against the crowd is beneficial. Alternatively, one could construct a game model derived from the *College Admission Problem* or *Stable Marriages Problem* (Gale and Shapley, 1962; Manlove and Sng, 2006; Abraham et al., 2007; Akbarpour et al., 2016). In these problems, agents try different matches until an optimal match is found. Yet, in vehicle for hire markets, I assume that redirecting to another resource, if the preferred resource is not available, is not an option, because of the costs and time constraints of redirecting (requires the agents to drive to another location consuming time and fuel).

To analyze the fundamental underlying problem, I propose a repeated non-cooperative game model to investigate different strategies in the coordination problem among drivers. It is a generalization of the Kolkata Paise Restaurant Problem (KPRP) (Chakrabarti et al., 2009) where agents repeatedly compete for a set of resources. As a foundation to be able to assess coordinators' incentives like 'surge pricing', one first needs to understand the fundamental impact of different driver strategies. I contribute to this research field by game model, relaxing assumptions of the KPRP. In contrast to existing research, I address both individual agent preferences and different resource capacities. Besides the game model, the contributions of this research are different mixed strategies for the model and an analysis of their impact on car utilization and driver utilities in different settings. These insights constitute building blocks for a characterization of favorable agent behavior to design incentive mechanisms to distribute drivers efficiently.

1.1. The Vehicle for Hire Problem and its Model Variants

In this thesis, I cover five different, but related model variants: The Kolkata Paise Restaurant Problem and four relaxations suited for mobility markets comprising the Vehicle for Hire Problem (VFHP).

In Kolkata, there were very cheap and fixed-rate 'Paise Restaurants', popular among the daily laborers in the city. During lunch hours, the laborers used to walk down (to save the transport costs) to any of these restaurants and would miss the lunch if they arrived at a restaurant where their number is more than the capacity of the restaurant for such cheap lunch. Walking down to the next restaurant would mean failing to report back to the job in time! Paise means the smallest Indian coin and there were indeed some well known rankings of these restaurants as some of them would offer more tastier items compared to the others. (Chakrabarti et al., 2009, p. 2421)

The KPRP was first presented by Chakrabarti et al. (2009). In this model, N agents (that is daily laborers) aim at having lunch at one of the N restaurants. All agents gain the same utility from some restaurant, and all restaurants have mutually different utilities. Every agent aims at getting lunch at his preferred restaurant, but every restaurant can only cater a single agent. Thus, if more than 1 agent goes to some restaurant, some agents will not get lunch, as they cannot divert to another restaurant that same day. The KPRP is a repeated game with an infinite number of iterations.

In mobility markets, drivers $i \in I$ constitute agents and customers $j \in J$ (located in districts $k \in K$) constitute resources. Agents drive to resources. Agents carry resources (up to the capacity limit). For this thesis I relax two main assumptions: Agents no longer retrieve identical utility from a given resource, but one agent can prefer resource j and another agent can prefer resource $j' \neq j$ (with the highest utility determining preference). I present two different models: In the Individual Preferences model (IP), utilities are uniformly assigned to resources (customers). Thus, agent preferences are independent of each other. In the Mixed Preferences model (MP), utilities are calculated as a weighted average of an individual component (that is distance between agent and customer) and a shared component (that is the payoff). I further model increased capacity: Clustering customers $j \in J$ in districts $k \in K$ allows agents to divert to other customers inside the district they drove to. The average number of customers per district is φ , and the customers randomly "choose" the district they belong to, the number of customers per district is thus Gaussian distributed around φ . The Individual Preferences with

Multiple Trips per Customer model (IPMC) combines the IP model with the clustering concept: Agents gain random utilities from customers and customers belong to districts. In the Mixed Preferences with Multiple Trips per Customer model (MPMC), the utility is obtained as a weighted average of an individual component to model the distance and a shared component to model the payoff. The distance (and thus the individual component) is equal for all customers belonging to one district.

1.2. Outline of this Thesis

The remainder of this thesis is organized as follows: I first discuss related work in chapter 2, I then present the strategies (chapter 3). The successive chapters present the individual model variants and assess the performance of aforementioned strategies. Chapter 4 focuses on the KPRP, chapter 5 presents the IP model variant, chapter 6 gives insight in the MP model variant, chapter 7 concerns the IPMC model variant, and chapter 8 evaluates the MPMC model variant. To improve the reader's understanding, chapters 4-8 can be read independently from each other, as key concepts are presented in each of them. Chapter 9 discusses the results from chapters 4-8, and chapter 10 concludes this thesis.

2. Related Work

To my knowledge, no paper extends the KPRP for mobility markets. Relevant research is conducted in three fields: First, I give an overview of relevant game models in other application areas, in particular coordination games. Second, there is literature in optimization and operations research in the field of vehicle for hire markets. Third, I introduce basic literature of dynamic mechanism design.

2.1. Congestion Games

The presented model is a type of congestion game, a model for games in which agents should choose different alternatives to succeed first described by Rosenthal (1973). Mathematically, congestion games can be identified by their potential function and thus their pure-strategy Nash-equilibria; Congestion games are therefore also Potential games (Monderer and Shapley, 1996; Nash, 1951). Yet, such a Nash equilibrium is usually inefficient, as Correa et al. (2005) prove. Other congestion game models are the El Farol Bar Problem (Arthur, 1994), the KPRP (Chakrabarti et al., 2009), the Crowding Game (Milchtaich, 1996), and the minority game (Challet and Zhang, 1998).

The El Farol Bar Problem is a game model with N agents (scientists) and one resource (the bar in Santa Fe during Karaoke night). All agents aim at maximizing their profit. If more than $0.6 \cdot N$ agents go to the bar, it becomes overcrowded, and the agents would enjoy themselves more at home. If fewer agents go to the bar,

they enjoy themselves more than if they stayed at home. Agents, therefore, coordinate themselves such that as many agents as possible (but less than $0.6 \cdot N$) go to the bar (Arthur, 1994).

The KPRP is the foundation game model for this thesis; the model is described in chapter 4 in more detail. Chakrabarti et al. (2009) and Ghosh et al. (2013) introduce strategies for increasing the utilization of the KPRP. Yang et al. (2016) study a generalization of the KPRP which is also aimed at dynamic markets: As a relaxation of the KPRP they study whether an agent should divert to another district or stay in the current one with different capacities for different districts. Agents are being replaced by others (which do not have the same prior knowledge) following a Poisson distribution. They prove the existence of a *Mean Field Equilibrium* (Lasry and Lions, 2007) for the *Threshold Strategy* (if a capacity threshold is exceeded at time t , agents stochastically divert to other districts) (Yang et al., 2016). This thesis on the opposite compares different strategies. Agarwal et al. (2016) generalize the KPRP to a *Majority Game*, in which they study convergence behavior given only few prior knowledge. In difference to the KPRP, capacity is not restricted, and in difference to the problem in mobility markets agents have no internal utility ranking, they aim at choosing with the herd.

The Crowding Game is a game model in which the utility of agents only depends on the number of agents also selecting the same option. If more agents select one option, the utility decreases (Milchtaich, 1996). The VFHP game model is similar to the Crowding Game as the number of agents decreases the utility (as the expected utility is divided among all agents selecting some resource), but this model also uses a basic utility which is not shared among agents.

The Minority Game is a game with N agents and two resources, and the utility for those agents choosing the resource with the lower occupancy is higher than the utility for those agents in the crowded resource (i.e. roads) (Challet and Zhang, 1998). In a recent study, "treatments" (which differ in the information given to participants) for the Minority Game were studied with experiments. The authors state that changing from one option to the other is not recommended regardless of prior knowledge (Chmura and Pitz, 2006). Because the Minority Game only allows two different payoffs from two different resources, I cannot directly transfer this insight to the Kolkata Paise Restaurant Problem in mobility markets.

2.2. Vehicle for Hire Market

There is only limited research work available on optimal distribution of drivers in vehicle for hire markets. Several studies focus on assigning drivers an optimal district where they await passengers (Lee et al., 2004; Seow et al., 2010); though, in most business models, drivers

decide independently. Yang et al. (2005) study a model with varying demand and supply. Taxi drivers individually decide when to enter the market and when to leave it, resulting in a market equilibrium. This work does not study utility, but only utilization. Kim et al. (2011) propose an agent-based model incorporating real-world passenger travel pattern to predict the highest possible utility. Their model also incorporates districts ("areas") and varying utility functions over time, but tests for different criteria: Whilst I analytically derive utilization and utility for different strategies in a large environment, Kim et al. (2011) studies a setting with five nodes and retrieves utilization and passenger wait time for varying fleet sizes. Wong's primary criterion is reduced vacant mileage for taxis (Wong et al., 2015). He uses a two-step approach in which taxis can only divert to adjacent zones rather than all others. Trigo et al. (2006) uses Multi-Agent Markov Decision Processes to model drivers transporting passengers. This paper uses a cover story which is highly similar to ours, but rather than using stochastic strategies, Trigo et al. (2006) use a two-layered learning process. This thesis aims at improving the taxi allocation with respect to utilization fraction or utility assuming choice at discrete time steps. Li (2006) on the opposite studies strategies to minimize passenger waiting time or travel time, taxi idle time or non-live mileage with drivers deciding asynchronously. This thesis studies a large variety of strategies, Li (2006) restricts himself to three simple strategies. The paper concludes that returning to hotspots after serving a trip can increase all studied parameters. Similar results can also be seen in this thesis, as the utilization fraction increases after introducing multiple trips per district.

Li et al. (2011) present a model which predicts whether agents should wait for passengers stationary or continue driving to "hunt down" customers. They use data mining techniques with data on time, location, and strategy (hunt or wait). In the VFHP, all agents decide where to drive to (yet, the location might not change). Thus, the strategy of the VFHP dictates where to go rather than if to go to another location. The model by Li et al. (2011) cannot predict where taxi drivers should drive. Ge et al. (2010) build a recommender system to reduce the travel distance before carrying the next customer. This behavior is reflected by the VFHP game model, as the individual utility models distance. Yuan et al. (2011) extend the work by Ge et al. (2010) by also recommending optimal passenger behavior.

Alonso-Mora et al. (2017) postulate that it should be possible to replace 13,000 cabs in New York City by only 3,000 on-demand vehicles for ride-sharing, which would both reduce wait time and traffic congestion. Their calculations suggest that a better utilization fraction of cabs can be achieved, though ride-sharing is not considered in this thesis. Furthermore, using graph traversals for

optimal distribution and routing of taxis is a solution a single driver cannot adopt, but only dispatchers.

Shi and Lian (2016) study the taxi transportation market from the opposite side as this thesis paper does: Passengers can decide whether or not they are queueing for a taxi (depending on the "queue length" (number of passengers) and the "buffer size" (number of cabs) at the taxi stand). The authors compare strategies of selfish and social passengers and options for the government to interfere.

Furthermore, there are several papers in the field of operations research which focus on the influence of regulation (taxi medallions, fixed rates) on the market (Cairns and Liston-Heyes, 1996; Arnott, 1996). In the VFHP game model, I assume that there are sufficient agents to carry every customer and sufficient customers such that every agent can carry a customer.

2.3. Dynamic Mechanism Design

There is early stage work on dynamic mechanism design in matching markets: If there is a dispatcher, he can make agents wait for a better suited trip. Kurino (2009) gives a dynamic version of the House Allocation Problem. Bloch and Houy (2012) periodically redistribute items between agents.

If agents are allowed to choose independently from a dispatcher, waiting time might influence their choice, reducing welfare. In this component – choosing the best individual option reduces social optimality – the problem described by Leshno (2012) is highly similar to the KPRP. Yet, unlike environments described in the paper (e.g., nursing homes, subsidized housing), there are no "overloaded waiting lists" (demand tremendously exceeds supply) in the taxi industry, as passengers usually have other means of transportation to choose from.

Social Welfare (benefit for the entire group) in transportation markets has been studied at the example of Rotterdam Port: Transportation tasks inside the port are assigned to trucks which are waiting for departure. The authors claim that a higher number of participants in general increases social welfare (as it is easier to adapt to peak load times), but agents might not continue participating if they assumed that the game put them at a disadvantage in comparison to other players. They, therefore, postulate an algorithm which ensures that agents are equally utilized (Ye and Zhang, 2016; Ye et al., 2017). In the KPRP on the opposite, I assume that the number of customers always equals the number of agents (agents will always participate), but agents are not assigned their trip.

Chen and Hu (2016) conduct research on market design in a market place with buyers and sellers such as Uber: In such markets, buyers wait for lower market prices while sellers wait for higher market prices. They conclude that fast changes in the market price (set by

an intermediary) and price surges are not recommended, as participants might leave the market temporarily. This thesis on the opposite assumes myopic agents, who only plan ahead few time steps.

3. Strategies

In this thesis I consider seven strategies: No Learning (NL), Rank Dependent Choice (RD), Limited Learning (LL), One Period Repetition (OPR), Crowd Avoiding (CA), Stochastic Crowd Avoiding (SCA), and Stochastic Rank Dependent Choice (SRD). NL and RD are baseline strategies which represent basic behavior. RD, LL, OPR, and SRD incorporate the resource's utility in the agents' choices and are therefore utility-based. LL, OPR, CA, and SCA require knowledge about previous iterations and are therefore history-based.

The NL strategy dictates agents to randomly choose a resource in every iteration, regardless of history (hence the term "No Learning") or resource utility. Resources are either customers or districts (or restaurants in the KPRP). The strategy was first presented by Chakrabarti et al. (2009) in which restaurants comprise resources.

The second baseline strategy is the strategy RD. Agents always drive to the resource yielding them maximum utility. Agents thus receive maximum utility, if they carry a customer. If there are several resources yielding equal utility, agents decide randomly between all maximum utility resources. I introduce this strategy, as it mimics simple behavior if limited information is available: If agents do not know about the preferences or behavior of other agents, but assume that only a few agents share the same preference, the most simple approach is to always head for the preferred resource. It requires only very few computational power: Prior to the first iteration, agents calculate their preferred customer by comparing the utility of all resources. After driving there, they will remain in their position, requiring no recomputation at all. It also requires no information except the own utilities or preferences, making it suitable for large problem spaces.

Agents incorporating the LL strategy follow a two-step approach: (1) If an agent carried a customer at time t , he will drive to the highest utility resource at time $t + 1$. (2) If an agent did not carry a customer at time t , he will randomly choose any other resource at time $t + 1$. (If an agent was successful at the highest utility resource, he will return there in the next iteration). The LL strategy was presented by Chakrabarti et al. (2009) (named Limited Learning 1).

The OPR strategy requires agents to follow a three-step approach: (1) If an agent carried customer j at time t (but not at time $t - 1$), he will return to this resource at time $t + 1$ (return). (2) If an agent served the same resource j at time $t - 1$ and t , he will compete for the highest utility

customer at time $t + 1$ (improve). (3) If an agent did not carry any customer at time t , he will randomly choose any resource which was vacant at time t in the next iteration (random). OPR was also introduced in Chakrabarti et al. (2009).

With the CA strategy agents only drive to resources which were vacant or had remaining capacity at time $t - 1$. This strategy originates in a paper by Ghosh et al. (2013).

Agents using the SCA strategy stochastically decide whether to return to the same resource or to randomly turn to another resource. If a resource j does not exceed its capacity at time t , all agents driving to this resource j at time t will return there at time $t + 1$. If the capacity is exceeded, all agents stochastically either return to j or drive to any other (randomly chosen) resource at time $t + 1$ such that the expected number of agents in j equals its capacity (let the capacity be c_j and the number of agents at the resource be o_j : return with probability $\frac{c_j}{o_j}$, randomly choose another resource with probability $1 - \frac{c_j}{o_j}$). The SCA strategy stems from Ghosh et al. (2013).

The SRD strategies build upon the RD strategy, including some properties of the SCA strategy: Let the capacity of a resource j be c_j , and let the number of agents preferring resource j be p_j (agents who cannot retrieve higher utility from any other resource). Agents drive to their preferred resource if its capacity is not exceeded, that is $c_j \geq p_j$. Otherwise, they stochastically drive to j with probability $\frac{c_j}{p_j}$ and redirect to another resource with probability $1 - \frac{c_j}{p_j}$. Thus, the expected number of agents preferring a resource j driving to that resource j is c_j , if at least c_j agent prefer j , and p_j otherwise. The resource agents divert to can be one of the following: (SRD1) Any customer which is noone's first choice; (SRD2) any other customer; (SRD3) his second choice customer; or (SRD4) the best customer which is noone's first choice. SRD3 and SRD4 are an extension of SRD2 and SRD1 respectively, increasing the average utility of successful agents, that is agents carrying a customer. If the first preferences of different agents are not independent, this likely also apply for the alternate preferences in SRD3 and SRD4, decreasing the utilization fraction. All SRD strategies require information about the first preferences of all other agents which can be acquired by a single iteration of RD upfront. Then all agents know how many other agents share the same top preference, making the second iteration identical (SRD2) or similar (SRD1, SRD3, SRD4) to SCA, as all agents redirect based upon the number of agents in the chosen district during the previous iteration. In addition to the number of agents preferring the same resource, the SRD1 strategy also requires information about the number of agents preferring all other resources which one could also retrieve in a single iteration of RD upfront.

Thus, the SRD1 strategy does not require too much information, if the number of iterations is sufficiently high to compensate for a potentially very low utility during the first iteration. The SRD2 strategy requires less information than the SRD1 strategy, as it only incorporates the number of agents preferring the resource they prefer themselves. It is thus beneficial if the information about other resources cannot be determined easily. The SRD3 strategy also requires only very few information (as much as SRD2), but the utility of successful agents is higher, as all successful agents receive a high utility (maximum or second highest utility). If the agent utilities of different agents are not stochastically independent, there can be a high number of resources noone drives to, neither as first nor as second preference. In many cases, the second preference of an agent is the first preference of another agent, thus not exploiting the full potential. In SRD4, the second preference is only chosen, if no agent prefers this resource. Thus, the set of first choice resources and the set of alternate choice resources do not intersect, making it impossible that alternate choice agents carry a customer who is preferred by another agent increasing the average utility. Yet, SRD4 requires more information about the preferences of other agents than SRD3. Thus, the existence of all strategies is justified by their different data requirements comparing to the expected performance. The performance of the different strategies with respect to the metrics utilization fraction and utility depends on the actual model variant.

4. Kolkata Paise Restaurant Problem

In their paper, [Chakrabarti et al. \(2009\)](#) discussed different strategies and provided simulations.

In the following, I will briefly reproduce their results analytically.

4.1. The Model

In the KPRP cover story, daily laborers $i \in I, |I| = N$ represent agents who select a restaurant $j \in J, |J| = N$ for lunch. Agents select (i.e. randomly) a restaurant to which they drive. Formally, I use $d(i, j)$ to represent that i goes to j .

$$d(i, j) = \begin{cases} 1 & \text{if agent } i \text{ goes to restaurant } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 4.1})$$

$$\forall j : o_j = \sum_{i \in I} d(i, j) \quad (\text{Definition 4.2})$$

Obviously, one agent can only go to one restaurant ($\forall i : \sum_{j \in J} d(i, j) = 1$). Every restaurant $j \in J$ can cater exactly one agent $i \in I$.

$$c(i, j) = \begin{cases} 1 & \text{if agent } i \text{ eats at restaurant } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 4.3})$$

If no agent went to j , j does not cater any agent, if more than one agent goes to restaurant j , only one will be served ($\forall j : c(i, j) = \min \left(\sum_{i \in I} d(i, j), 1 \right)$). Agents can only eat at restaurants they went to ($\forall i, j : c(i, j) \leq d(i, j)$). The utility $u(i, j)$ agents receive from eating at a restaurant is a random permutation and is identical for all agents (resulting in a shared utility $u_s(j)$), that is $\forall j : u(i, j) = u_s(j)$ and $\forall j, j' : u_s(j) \neq u_s(j') \vee j = j'$. A daily laborer (agent) prefers a restaurant if no other restaurant yields higher utility for him. The number of agents preferring a restaurant j is denoted as p_j .

$$p(i, j) = \begin{cases} 1 & \text{if } \forall j' \in J \setminus \{j\} : u(i, j) \geq u(i, j') \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 4.4})$$

$$\forall j : p_j = \sum_{i \in I} p(i, j) \quad (\text{Definition 4.5})$$

The utilization fraction f is given as the number of agents getting lunch divided by the total number of agents. If an agent i gets lunch is given by $f(i)$ which is 0, if i ate at no restaurant, and 1 otherwise (as every agent can eat at maximum one restaurant).

$$f = \frac{1}{N} \cdot \sum_{i \in I} f(i) \quad (\text{Definition 4.6})$$

$$f(i) = \sum_{j \in J} c(i, j) \quad (\text{Definition 4.7})$$

The overall utility u is average utility per agent. The agent utility $u(i)$ is $u(i, j)$, if i eats at j and 0 otherwise.

$$u = \frac{1}{N} \cdot \sum_{i \in I} u(i) \quad (\text{Definition 4.8})$$

$$u(i) = \sum_{j \in J} u(i, j) \cdot c(i, j) \quad (\text{Definition 4.9})$$

In experiments and simulations, I further assume $N = 1000$ (1000 agents and 1000 customers), and that customers are indexed by their utility ($u_s(j) = \frac{j}{N}$). Thus, the utility is uniformly distributed such that $u_m = u_{max} = 1$ is the utility of agents eating at their preferred restaurant, and $u_{avg} = 0.5$ is the expected utility of agents eating at any other restaurant.

4.2. Theoretic Foundations

The capacity of all restaurants is 1. All agents prefer the same restaurant j_p . Thus, the probability that $j \in J$ is preferred by exactly p_j agents is 1 for j_p and $p_j = N$ and 0 otherwise.

4.3. No Learning

As a baseline comparison Chakrabarti et al. (2009) give an entirely random selection: In every iteration, every agent selects one of the restaurants at random.

In Chakrabarti et al. (2009) they give the formula equation 1 as probability $P(o_j)$ for o_j agents choosing the same restaurant, if on average λ agents go to the same restaurant. Equation 2 simplifies equation 1 by setting $\lambda = 1$. With $N \rightarrow \infty$, one can further simplify the formula using the Poisson Limit Theorem.

$$P(o_j) = \binom{\lambda N}{o_j} \frac{1}{N} \left(1 - \frac{1}{N}\right)^{\lambda N - o_j} = \frac{\lambda^{o_j}}{o_j!} e^{-\lambda} \quad (1)$$

$$= \binom{N}{o_j} \left(\frac{1}{N}\right)^{o_j} \left(1 - \frac{1}{N}\right)^{N - o_j} = \frac{1}{o_j!} e^{-1} \quad (2)$$

Therefore, $P(0)$ gives the probability of a restaurant being unoccupied any evening using this random strategy, making $1 - P(0) \approx 63.2\%$ the average utilization.

I, therefore, expect a Gaussian distribution around $f = f_{NL} = 63.2\%$ for the utilization fraction. As agents on average receive average utility (if they are successful), I conclude that the utility is $u = f \cdot u_{avg} = 0.316 \cdot u_{max}$.

4.4. Rank Dependent Choice

Agents $i \in I$ incorporating the RD strategy always turn to the restaurant j that yields them the highest utility ($d(i, j) = 1 \iff \forall j' : u(i, j) \geq u(i, j')$).

In the KPRP, the restaurant with the highest utility and thus the first preference restaurant is identical for all agents ($\forall i, i' \in I : u(i, j) = u(i', j)$). Thus, all agents $i \in I$ go to the same restaurant j . This restaurant can only cater a single agent, resulting in a utilization fraction of $f = \frac{1}{N}$. For $N = 1000$, I, therefore, expect $f = f_{RD} = 0.1\%$. The (single) successful agent receives maximum utility, resulting in $u = 0.001 \cdot u_{max}$ on average.

4.5. Limited Learning

With this strategy, all agents choose a restaurant at random the first night. The utilization therefore is Gaussian distributed around 63.2%. During successive nights, all agents base their choice on whether they got dinner the previous day (Chakrabarti et al., 2009):

- If some agent got food at time t , he will choose the highest ranking restaurant at time $t + 1$. (If an agent was successful at the highest utility restaurant, he will return there in the next iteration)

- If some agent did not get food at time t , he will randomly choose any other restaurant at time $t + 1$.

The first case is irrelevant for the KPRP, as the utilization fraction for this part is $f_{RD} = \frac{1}{N}$ (with f_{RD} as the utilization fraction of the RD strategy or fraction of carried customers by an agent preferring them), with $N \rightarrow \infty$ the utilization fraction gets negligibly small (or $f_{RD} = 0.1\%$ for $N = 1000$). The second case is given by $\lambda = 1 - f$ in equation 1 (the ratio between agents and restaurants is $(1 - f) : 1$). Chakrabarti et al. (2009) give the following recursion relation:

$$f_{t+1} = 1 - e^{-\lambda t}; \lambda_t = 1 - f_t \quad (3)$$

In a more generalized fashion, I write:

$$f_t = \underbrace{f_{t-1} \cdot f_{RD}}_{\text{first try best}} + \underbrace{\left(1 - e^{-(1-f_{t-1})}\right)}_{\text{random or return}} \quad (4)$$

If one assumes that f converges as $f_{t+1} = f_t$, the utilization will be Gaussian distributed around an average value of $f = 43.3\%$ and $u = f \cdot u_{avg} = 0.212 \cdot u_{max}$.

4.6. One Period Repetition

All agents choose the restaurant randomly the first evening.

- If some agent got dinner at restaurant j at time t (but not at time $t - 1$), he will return to this restaurant at time $t + 1$ (*return*).
- If some agent got dinner at the same restaurant j at time $t - 1$ and t , he will compete for the highest utility restaurant at time $t + 1$ (*improve*).
- If some agent did not get dinner at any restaurant at time t , at time $t + 1$ he will randomly choose any restaurant which was vacant at time t (*random*).

In their paper, Chakrabarti et al. (2009) both give the distribution and simulation results.

The probability distribution of utilizations is given by equation 6 with x_t being the fraction of agents returning to the same restaurant at time $t + 1$, and thus the fraction of agents eating at a randomly chosen restaurant at time t . As all agents who do not eat at a restaurant at time $t - 1$ choose a restaurant randomly and are successful with probability f_{NL} , Chakrabarti et al. (2009) assume that $x_t = (1 - x_{t-1}) \cdot f_{NL}$. x_t is also the fraction of agents improving at $t + 2$ (in this case, the expected utilization is $f_{RD} = \frac{1}{x_t N}$, it can therefore be ignored if $N \rightarrow \infty$).

$$f_t = x_{t-1} + (1 - x_{t-1}) \cdot \left(1 - e^{-1}\right) \quad (5)$$

$$f_{t+1} = (1 - x_t) \cdot (1 - e^{-1}) + \left(1 - (1 - x_t) \cdot (1 - e^{-1})\right) \cdot (1 - e^{-1}) \quad (6)$$

In their paper, Chakrabarti et al. (2009) conclude that the fixed point of this right half of both equations in 6 is at $x \approx 0.38$ or $f \approx 0.77$, a result I cannot replicate in simulations.

Their original formula is not replicable: It only considers those agents who are not eating at their preferred restaurant (the utilization fraction for these agents is added in the second term). From the remaining $(1 - f_{RD}) \cdot N$ agents, a fraction of x_{t-1} agents returns to the previously chosen restaurant, and a fraction x_{t-2} tries eating at the highest utility restaurant (yet unsuccessful, as all successful agents contribute utilization via the second term). Thus, a fraction of $(1 - x_{t-1} - x_{t-2})$ of all agents randomly chooses a restaurant. These agents are successful with probability $f_{NL} = 1 - e^{-1}$. Chakrabarti et al. (2009) do not deduct x_{t-2} , as these agents are unsuccessful. In the next iteration, those agents who successfully randomly choose a restaurant $((1 - x_{t-1} - x_{t-2}) \cdot f_{NL})$, become x_t . Assuming that x_t converges to a stable state ($x_t = x_{t+1} = x_{t+2}$), I can drop subscript t , resulting in a fraction x . The corrected formula is given in equation 7.

$$f = \underbrace{(x + (1 - 2 \cdot x) \cdot f_{NL}) \cdot (1 - f_{RD})}_{\text{random, return, and improve}} + \underbrace{f_{RD}}_{\text{best}} \quad (7)$$

The fraction x is given by $x = (1 - 2x) \cdot (1 - e^{-1}) \approx 27.9\%$, and f_t decreases to $f = 55.8\%$. The utility is given as $u = 0.279 \cdot u_{max}$.

Yet, one should notice that this strategy is promising for vehicle for hire markets: The best (highest utility) resources are different for different agents, thus, this share is not "lost", but will be added.

4.7. Crowd Avoiding

Agents using the CA strategy only choose restaurants which did not serve customers the previous evening.

The probability $P(0)$ of a restaurant being vacant at time $t = 1$ after being empty at time $t = 0$ is given by equation 8. As the number of restaurants to choose from at time $t = 1$ is reduced from 1 to $1 - f$, the average number of agents per restaurant needs to be set to $\lambda = \frac{1}{1-f}$ to cater for this change (in equation 1).

$$P(0) = e^{-\lambda} = e^{-\frac{1}{1-f}} \quad (8)$$

Incorporating $f = 1 - P(0)$ and the fact that only $1 - f$ restaurants are available into equation 8, yields the following equation:

$$f = (1 - f) \left(1 - e^{-\frac{1}{1-f}}\right) \quad (9)$$

Equation 9 has two solutions at $f_1 \approx 0.457$ and $f_2 \approx 1.872$, the latter being discarded as the utilization fraction cannot exceed 1. The utilization fraction is therefore $f = 45.7\%$. As all agents who eat at any restaurant receive average utility, I conclude that the utility is $u = 0.229 \cdot u_{max}$.

4.8. Stochastic Crowd Avoiding

Ghosh et al. (2013) also introduced another strategy in which the probability of returning to some place inversely depends on the number of agents choosing this restaurant ($ret_j(t) = \frac{1}{o_j(t-1)}$ with ret_j the probability of returning to restaurant j and $o_j(t-1)$ the number of agents at restaurant j at time $t-1$). Alternatively, this agent will choose any other restaurant with equal probability $\frac{o_j(t-1)-1}{o_j(t-1)} \cdot \frac{1}{N-1}$.

In their paper, Ghosh et al. (2013) give an expected utilization fraction of $f \approx 80\%$. My simulations give an average utilization fraction of $\bar{f} = 0.735$. This is still better than random (the only better than average strategy), but it is not as good as expected.

Ghosh et al. (2013) define that a_i is the share of restaurants with i agents (in our model, i is o_j) and $a_i = 0 \forall i > 2$. Thus, $a_0 + a_1 + a_2 = 1$ (number of restaurants), and $a_1 + 2 \cdot a_2$ (number of agents). In every iteration, the share of vacant restaurants (a_0) is newly calculated, it comprises those restaurants which were empty the previous iteration (prev), minus those restaurants to which some agent drives to who went to an a_2 restaurant the previous iteration (new) and those a_2 restaurants in which both agents from the previous iteration divert and no agent goes to (both leave).

$$a_0 = \underbrace{a_0}_{\text{prev}} - \underbrace{a_0 \cdot a_2}_{\text{new}} + \underbrace{\frac{a_2}{4} - a_2 \frac{a_2}{4}}_{\text{both leave}} \quad (10)$$

I assume that the difference emerges from the fact that the authors ignored that more than two agents can head for in the same restaurant. They state that the influence of a_i for $i > 2$ is negligibly small), yet, using $a_0 = a_2 + 2 \cdot a_3 + 3 \cdot a_4 + \dots$ the accumulated impact grows. In simulations with $N = 1000$ agents, I observed $o_j = 3$ in 3.39% of all restaurants and $o_j = 4$ in 0.42% of all restaurants, $o_j = 5$ to $o_j = 10$ occurred seldom, but still affected the final result.

The utility is $u = f \cdot u_{avg} = 0.368 \cdot u_{max}$.

4.9. Stochastic Rank Dependent Choice

Agents using the SRD strategy stochastically either eat at the highest-utility restaurant j_p or turn to another restaurant $j \in J$. As all N agents share the same first preference, the probability that some agent i goes to j_p is $\frac{1}{N}$.

In the SRD1 strategy, the other agents turn to all restaurants except j_p . On average, $N - 1$ agents turn to $N - 1$ restaurants, yielding an average utilization fraction of $1 - e^{-1}$ (for those $N - 1$ diverting agents). The total utilization fraction is therefore $f = \frac{1}{N} + \frac{N-1}{N} \cdot (1 - e^{-1}) = 63.2\%$ and the utility is $u = 0.316 \cdot u_{max}$.

In the SRD2 strategy, redirecting agents turn to all restaurants $j \in J$ (including j_p). On average, $N - 1$ agents turn to N restaurants, with $N \rightarrow \infty$ this yields an average utilization fraction of $1 - e^{-1}$ for diverting agents and an overall utilization fraction of $f = 63.2\%$ and a utility of $u = 0.316 \cdot u_{max}$.

In the SRD3 strategy, diverting agents turn to their second choice (that is the restaurant yielding second highest utility). As all utilities are identical for all agents, this second preference is shared among all agents. Thus, all diverting agents go to the same restaurant j' , resulting in a total utilization fraction of $f = \frac{2}{N} = 0.2\%$ for $N = 1000$ and a utility of $u = 0.002 \cdot u_{max}$.

The SRD4 strategy is identical to the SRD3 strategy for the KPRP, as the best vacant restaurant assuming all agents prefer the same restaurant is the restaurant that yields the second highest utility. I, therefore, conclude that the utilization fraction is $f = \frac{2}{N} = 0.2\%$ for $N = 1000$ and that the utility is $u = 0.002 \cdot u_{max}$.

4.10. Results

Table 1 comprises analytical and simulation results of the previous sections (simulation for SCA, analytical otherwise).

For the KPRP, utilization fraction and utility are linearly dependent for most strategies ($u = f \cdot u_{avg}$). RD, SRD3 and SRD4 have $u = f \cdot u_{max}$, but the performance with respect to utilization fraction or utility of these strategies is insufficient. All strategies exceed the baseline comparison RD, but only SCA outperforms the baseline NL. SRD1 and SRD2 are as good as NL, but cannot outperform it. SRD1 and SRD2 as well as SRD3 and SRD4 perform pairwise equally well, as the alternate choice is identical for the KPRP.

5. Individual Preferences

In this chapter, I will apply the strategies introduced in chapter 3 to the IP model variant. Some of the aforementioned strategies do not draw upon the actual ranking; I can therefore safely assume that the utilization will be the same as in the KPRP with the given adjustments.

5.1. The Model

I formally define the IP game as follows:

The utility agents $i \in I, |I| = N$ receive from carrying some customer $j \in J, |J| = N$ is uniformly distributed, that is every agent associates every utility level between 0 and 1 with $\frac{1}{N}$ step size with some customer, but different agents may receive different payoff from the same customer. I assume strict utility levels (no two customers are associated with the same utility by some agent) and are therefore able to derive a preference ranking.

Every agent $i \in I$ drive to exactly one customer $j \in J$ ($\forall i : \sum_{j \in J} d(i, j) = 1$). I denote that i drives to j as $d(i, j) = 1$.

The number of agents driving to some customer j is its occupancy o_j .

$$d(i, j) = \begin{cases} 1 & \text{if } i \text{ drives to } j, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{Definition 5.1})$$

$$\forall j : o_j = \sum_{i \in I} d(i, j) \quad (\text{Definition 5.2})$$

Every agent drives to exactly one customer ($\forall i : \sum_{j \in J} d(i, j) = 1$). If more than one agent drives to some customer j , only one of the agents will be able to carry j ; all others will run empty. I denote that agent i carries customer j as $c(i, j) = 1$.

$$c(i, j) = \begin{cases} 1 & \text{if } i \text{ carries } j, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{Definition 5.3})$$

Obviously, an agent i can only carry a customer j , if he drives to j ($\forall i, j : c(i, j) \leq d(i, j)$). A customer j is carried by at most one agent, and if there is an agent i that drives to j , this customer will be carried ($\forall j : c(i, j) = \min\left(\sum_{i \in I} d(i, j), 1\right)$). Agents can either randomly or deterministically choose the customer they drive to. Every agent prefers one customer over all others, as it returns the highest utility for him (if no other agents were driving to the same customer). This customer j yields a higher utility than all other agents. The number of agents preferring some customer j is denoted as p_j .

$$p(i, j) = \begin{cases} 1 & \text{if } \forall j' \in J \setminus \{j\} : u(i, j) \geq u(i, j') \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 5.4})$$

$$\forall j : p_j = \sum_{i \in I} p(i, j) \quad (\text{Definition 5.5})$$

The utilization fraction is derived from the average number of agents carrying a customer.

Strategy	utilization f	utility u
NL	63.2%	0.316
RD	0.1%	0.001
LL	43.3%	0.212
OPR	55.8%	0.279
CA	45.7%	0.229
SCA	73.5%	0.368
SRD1	63.2%	0.316
SRD2	63.2%	0.316
SRD3	0.2%	0.002
SRD4	0.2%	0.002

Table 1: KPRP: Comparing Strategies

$$f = \frac{1}{N} \cdot \sum_{i \in I} f(i) \quad (\text{Definition 5.6})$$

$$f(i) = \sum_{j \in J} c(i, j) \quad (\text{Definition 5.7})$$

The utility is given as the average utility of all agents. The individual utility $u(i, j)$ an agent i receives from carrying a customer j is a random permutation for every customer ($\forall i : \forall j, j' : u(i, j) \neq u(i, j') \vee j = j'$).

$$u = \frac{1}{N} \cdot \sum_{i \in I} u(i) \quad (\text{Definition 5.8})$$

$$u(i) = \sum_{j \in J} u(i, j) \cdot c(i, j) \quad (\text{Definition 5.9})$$

In numerical experiments and simulations I use $|I| = |J| = N = 1000$ agents and customers, and a uniformly distributed utility (between $\frac{1}{N} \approx 0$ and $u_{max} = 1$). Agents i carrying their preferred customer ($\forall j : c(i, j) = p(i, j)$) receive an expected maximum utility $u_m = u_{max}$, agents carrying another (not preferred) customer ($\sum_{j \in J} c(i, j) = 1 \wedge \forall j : p(i, j) = 1 \Rightarrow c(i, j) = 0$) receive u_{avg} .

5.2. Theoretic Foundations

The capacity of all customers is 1. The agent preferences are randomly distributed, Thus, the probability that p_j agents prefer customer j is Poisson distributed around 1.

$$Pref(p_j) = \frac{1}{p_j!} \cdot e^{-1} \quad (11)$$

5.3. No Learning

One of the best strategies for the *Kolkata Paise Restaurant Problem* with respect to the utilization fraction was

to choose a restaurant randomly at every evening. I will therefore adopt this strategy for mobility markets.

With this strategy, every driver randomly selects the customer (independent of his individual preference ranking and the history). Thus, the utilization fraction is calculated as $f = 1 - e^{-1}$ and is therefore $f = f_{NL} = 63.2\%$.

As agents choose randomly, on average every driver can expect utility u_{avg} . As only 63.2% of all drivers can expect payoff (the others do not get a customer), only those can get payoff. The average utility is therefore given by equation 12. In the given experiment with $N = 1000$ agents, I, therefore, expect a Gaussian distributed utility around an average of $u = 0.316 \cdot u_{max}$.

$$u = u_{avg} \cdot f = u_{avg} \cdot (1 - e^{-1}) \quad (12)$$

5.4. Rank Dependent Choice

The RD strategy is a second baseline comparison in addition to the NL strategy. Whilst the RD strategy was outperformed with respect to both metrics by all other strategies in the KPRP, the high number of distinct first preference resources makes it a reasonable choice in the IP model variant.

Assuming a random preference ranking, it would be beneficial to always try to get the maximum payoff, which – on average – should also yield an average utilization of $f = f_{RD} = 63.2\% = 1 - Pref(0) = 1 - e^{-1}$ with $Pref(0)$ being the probability that a customer is noone's first choice ($p_j = 0$). The expected average utility for successful agents – that is agents carrying a customer – increases from u_{avg} to u_{max} . In our example, this would be $u = 0.632 \cdot u_{max}$.

5.5. Limited Learning

Using the LL strategy, agents choose a customer randomly at time t and go to their highest utility customer at time $t + 1$, if they got a tour at time t , otherwise they choose randomly again.

The utilization fraction can be given by the following formula:

$$f_t = \underbrace{f_{t-1} \cdot f_{RD}}_{\text{first try best}} + \underbrace{\left(1 - e^{-(1-f_{t-1})}\right)}_{\text{random or return}} \quad (13)$$

The left summand of the equation models all those agents which chose their top priority customer at time t after successfully choosing randomly (at time $t - 1$ or earlier). The success rate for these agents is f_{RD} which is the utilization fraction of the RD strategy. The second summand of the equation comprises all those agents which choose randomly or which successfully chose their top priority at time $t - 1$ and return there. Using this equation, the utilization fraction is $f = 70.2\%$.

One has to differentiate between those agents who return to their prioritized customer and those agents who randomly choose a customer, as both belong to the second summand of equation 13. Let's assume that all those agents who do not share their top priority with any other agent will be able to return there. The fraction of returning agents is, therefore, given as $r = Pref(1) = e^{-1}$ (probability that $p_j = 1$ agents prefer a customer j).

The utility is given by 14 which results in a utility of $u = 0.620 \cdot u_{max}$ for $N = 1000$ for the IP model.

$$u = f \cdot f_{RD} \cdot u_m + \left(1 - e^{f-1}\right) \cdot \left(r \cdot u_m + (1 - r) \cdot u_{avg}\right) \quad (14)$$

5.6. One Period Repetition

Though the average utilization fraction was quite low for the One Period Repetition strategy in the KPRP, it can be a good solution for mobility markets: In the KPRP with identical rankings, the fraction of agents which headed for the best possible resource was usually lost (only one of them got dinner). This does not happen in mobility markets, as agents turn to different customers when going to their preferred resource.

Drawing upon the conclusions for the One Period Repetition in equation 7, I can assume that the new average utilization fraction is given by equation 16. Over time, all customers who are someone's first preference will be carried (second summand). f_{RD} is the utilization fraction of the RD strategy and, therefore, the fraction of customers carried by an agent preferring them. All other customers ($1 - f_{RD} = e^{-1}$) will be serviced during the random step and the improve step.

$$f = \left(x + (1 - 2x) \left(1 - e^{-1}\right)\right) \cdot (1 - f_{RD}) + f_{RD} \quad (15)$$

$$= \left(x + (1 - 2x) \left(1 - e^{-1}\right)\right) \cdot e^{-1} + \left(1 - e^{-1}\right) \quad (16)$$

Solving equation 16 yields an average utilization fraction $f = 83.7\%$.

The average utility is given by equation 5.6, in this formula, all those customers who are some agent's first preference will be serviced with maximum utility and all others will be serviced resulting in average utility for the respective agent. The result for this equation is $u = 0.728 \cdot u_{max}$.

$$u = \left(x + (1 - 2x) \left(1 - e^{-1}\right)\right) \cdot e^{-1} \cdot u_{avg} + \left(1 - e^{-1}\right) \cdot u_m \quad (17)$$

5.7. Crowd Avoiding

The strategy CA is identical to the one given in section 4.7 for the KPRP: All agents go to customers $j \in J$ who were vacant the previous iteration ($o_j = 0$ at time $t - 1$).

As this is strategy is independent of the rank, the expected utilization fraction is $f = 45.7\%$ from equation 9, and the utility is $u = f \cdot u_{avg} = 0.229 \cdot u_{max}$ for $N = 1000$ agents.

5.8. Stochastic Crowd Avoiding

Like in the CA strategy (section 5.7), the strategy SCA for mobility markets works exactly like the one for the KPRP in section 4.8: The probability of returning to a customer the successive day is inversely dependent on the number of agents at this customer the previous day.

This strategy is also independent of the actual utility resulting in expected utilization fraction of $\bar{f} = 0.735$ and a utility of $u = f \cdot u_{avg} = 0.368 \cdot u_{max}$ for $N = 1000$ agents.

5.9. Stochastic Rank Dependent Choice

Assuming every agent knows the number of agents p_j with an identical highest-ranking customer, agents could head for this customer with a probability of $\frac{1}{p_j}$ and head for either

- any customer which is noone's first choice (SRD1)
- any other customer (SRD2)
- his second choice customer (SRD3)
- the best customer which is noone's first choice (SRD4)

with a probability of $1 - \frac{1}{p_j}$.

The expected utilization fraction f is the sum over the utilization given p_j agents preferring some customer j for all possible values of p_j . $F(p_j)$ is the expected

fraction of customers being carried both in this customer and by switching to another customer (a more detailed description will follow in this section). $Pref(p_j)$ is the probability that some customer is preferred by p_j agents and is given by equation 11.

$$f = \sum_{p_j=1}^N Pref(p_j) \cdot F(p_j) \quad (18)$$

The fraction of agents servicing a customer given the number of agents preferring this customer p_j depends on the number of agents r_j switching (“redirecting”) to another customer. Every r_j is associated with a probability $D(p_j, r_j)$ that r_j out of p_j agents divert to other customers. Every agent that switches to another customer yields utilization with probability s (success rate). In total r'_j/r''_j agents receive this payoff. If at least one agent remains at this prioritized customer, this agent (or one of these agents) i will receive utilization $f(i) = 1$. (In SRD2 and SRD3 it is possible that redirecting agents turn to a customer in which at least one agent remains. In this case, diverting agents can “bully out” other agents. This is included in the success rate s .)

$$F(p_j) = \sum_{r'_j=1}^{p_j} D(p_j, r'_j) \cdot s \cdot r'_j + \sum_{r''_j=0}^{p_j-1} D(p_j, r''_j) \quad (19)$$

The probability that r_j out of p_j agents redirect to another customer is given by $D(p_j, r_j)$. Agents service their top priority customer with $p = \frac{1}{p_j}$, otherwise they redirect. For larger r_j and p_j , one can apply the Poisson Limit Theorem.

$$D(p_j, r_j) = \binom{p_j}{r_j} \left(\frac{1}{p_j}\right)^{p_j-r_j} \left(1 - \frac{1}{p_j}\right)^{r_j} \quad (20)$$

$$= \frac{1}{(p_j - r_j)!} \cdot e^{-1} \quad (21)$$

The average utility is given by adapting equation 18. The utilization fraction for p_j agents preferring the same customer is replaced by the utility $U(p_j)$ which gives the corresponding utility.

$$u = \sum_{p_j=1}^N Pref(p_j) \cdot U(p_j) \quad (22)$$

$U(p_j)$ modifies $F(p_j)$ by introducing different expected utilities for successful agents: If an agent switches to another customer, he can only expect average utility

u_{alt} , whilst staying with the top priority yields optimal utility u_m .

$$U(p_j) = \sum_{r'_j=0}^{p_j-1} D(p_j, r'_j) \cdot s \cdot r'_j \cdot u_{alt} + \sum_{r''_j=1}^{p_j} D(p_j, r''_j) \cdot u_m \quad (23)$$

The success rate s and the utilities u_m and u_{alt} depend on the behaviour of diverting agents. Table 2 lists these parameters, and they are discussed in the following sections.

5.9.1. Noone’s First Choice (SRD1)

The success rate s is given by on average e^{-1} agents switching over to other (vacant) customers. On average, e^{-1} customers are vacant.

$$s = (1 - e^{-1}) \quad (24)$$

I, therefore, derive $f = 79.5\%$ and $u = 0.678 \cdot u_{max}$.

I further assume $u_m = u_{max} = 1$ and $u_{alt} = u_{avg} = 0.5$, as agents redirect to a randomly selected customer.

5.9.2. Any Other Customer (SRD2)

The success rate s for redirecting agents changes in comparison to the previous strategy: If an agent frequents a customer who is someone else’s first preference, I cannot assume that the utilization is increased. On average, $e^{-1} \cdot N$ agents divert to other customers, and there are N customers these agents can divert to. The success rate is the probability that a diverting agent carries a customer j who is not preferred by any other agent ($p_j = 0$). On average, e^{-1} customers are not preferred by any agent. The probability that a customer j with $p_j = 0$ is not carried by another diverting agent is $e^{-\lambda}$ with λ the average number of diverting agents driving to a customer ($\lambda = \frac{1}{e^{-1}}$). Thus, the probability that at least one agent drives to some customer j is $1 - e^{-\frac{1}{e^{-1}}}$. The success rate is, therefore, given by equation 25.

$$s = e^{-1} \cdot \left(1 - e^{-\frac{1}{e^{-1}}}\right) \approx 0.347 \quad (25)$$

The expected maximum utility u_m is derived from the probability that $a = p_j - r_j$ agents remain with their shared first priority customer and another b agents get to this customer when selecting any other but their preferred customer. a agents remain if $r_j = p_j - a$ agents divert which is given by $D(p_j, p_j - a)$ from equation 21. The probability that b agents choose this customer randomly is given by equation 27 (swap to customer j). On average,

Strategy	s	u_m	u_{alt}	f	u
SRD1	0.623	1.00	0.50	79.5%	0.678
SRD2	0.347	1.00	0.50	69.0%	0.626
SRD3	0.347	1.00	1.00	69.0%	0.690
SRD4	0.623	1.00	1.00	79.5%	0.795

Table 2: IP: SRD Choice Strategy – Variables

e^{-1} of all agents divert to another customer, they choose from all N customers, thus, $\lambda = \frac{1}{e^{-1}} = e$.

$$u_m = \sum_{a=1}^N \sum_{b=1}^N D(p_j, p_j - a) \cdot P(b \text{ swap}) \cdot \frac{a \cdot u_{max} + b \cdot u_{alt}}{a + b} = 0.922 \quad (26)$$

$$P(b \text{ swap}) = \frac{\lambda^b}{b!} \cdot e^{-\lambda}, \lambda = e \quad (27)$$

The utilization fraction is, therefore, given as $f = 69.0\%$ and the utility is $u = 0.626 \cdot u_{max}$.

5.9.3. Second Choice Customer (SRD3)

In this strategy, every agent who knows that other agents share the same #1 priority decides to go to his #2 priority with probability $\frac{p_j-1}{p_j}$ (with p_j from Definition 5.5).

Success rate $s = 0.347$ and expected utility for successful non-diverting agents $u_m = 1.0$ remain unchanged with respect to SRD2, but u_{alt} for successful diverting agents increases to u_m . In the numerical experiment, the top priority customer yields a utility of 1.0, the second best had a utility of 0.999. Thus, the payoff is always either 1 oder 0.999 (And, therefore, $f \cdot 0.999 < u < f \cdot 1$). With $N \rightarrow \infty$ I can assume $u_{alt} = u_{max}$.

The utilization fraction is $f = 69.0\%$ and the utility is $u = 0.690 \cdot u_{max}$.

5.9.4. Best Vacant Customer (SRD4)

Rather than choosing any vacant customer (like in the first case), or always the second best (regardless of other agents choosing this customer as #1) an agent chooses the best possible customer in which no other agent might be serving with maximum utility.

Mathematically, choosing this alternative customer is identical to randomly choosing any vacant customer (there are $e^{-1} \cdot N$ vacant customers, as the customers are assigned as a random permutation, one could also randomly draw these customers). Therefore, the success rate of diverting agents is $s = 1 - e^{-1}$ like in SRD1 (equation 24). The utilization fraction is, therefore, $f = 79.5\%$.

If an agent approaches his top priority customer and is the only one there, the utility will be given by u_{max} . If the agent diverts to another customer, the expected utility is slightly lower. The highest utility customer cannot be the best vacant customer. The second best customer is vacant with probability e^{-1} . The customer with the third highest utility is vacant with probability e^{-1} , but only is the best vacant customer, if the customer with the second highest utility is not vacant (with probability $1 - e^{-1}$). The l best customer is the best vacant customer if all $l - 2$ customers (all customers yielding a higher utility except the first preference customer) are not vacant and customer l is vacant. Customer l then yields a utility of $1 - \frac{l}{N}$.

$$u_{alt} = \sum_{l=2}^N \left(1 - \frac{l}{N}\right) \cdot e^{-1} \cdot \left(1 - e^{-1}\right)^{l-2} = u_{max} - \frac{1.7183}{N} \quad (28)$$

For $N = 1000$, the utility of the alternate choice is $u_{alt} = 99.8\%$. With increasing N , this deviation becomes negligible ($u_{alt} \approx u_{max}$). The utility is, therefore, $u = 0.795 \cdot u_{max}$.

5.10. Results

Table 3 lists utilization fraction and utility for all strategies in this setting.

The two baseline comparison strategies NL and RD perform equally well with respect to f , but RD outperforms NL by orders of 2 concerning u , as all successful agents receive u_m (utility for agents carrying their preferred customer) rather than u_{avg} (average utility for agents carrying any customer). Except for CA, all strategies outperform NL and RD with respect to f (and NL with respect to u), but LL, SCA, and SRD2 fall behind RD with respect to utility, as agents receive a lower utility if they are successful (due to the fact that agents frequently choose a random customer). OPR performs best with respect to utilization but is outperformed by SRD4 regarding the utility. SRD1 and SRD4 as well as SRD2 and SRD3 show equal utilization, as the success rate is identical, but SRD3 and SRD4 outperform their counterparts on utility, as all agents receive (almost) u_{max} .

Strategy	utilization f	utility u
NL	63.2%	0.316
RD	63.2%	0.632
LL	70.2%	0.602
OPR	83.7%	0.728
CA	45.7%	0.229
SCA	73.5%	0.368
SRD1	79.5%	0.678
SRD2	69.0%	0.626
SRD3	69.0%	0.690
SRD4	79.5%	0.795

Table 3: IP: Comparing Strategies

6. Mixed Preferences

This section evaluates the performance regarding utilization and utility for the strategies defined in chapter 3 for the MP model: The distance to a customer is modeled as individual component in the utility of a customer, the payoff is modeled as the shared component.

6.1. The Model

The MP game is defined as follows: Agents $i \in I, s.t. |I| = N$ drive to customers $j \in J, s.t. |J| = N$ ($d(i, j) = 1$), agents try to carry the customer they drive to ($c(i, j) = 1$), but one customer can only be carried by one agent ($\forall j : c(i, j) = \min_{i \in I} \sum d(i, j), 1$). Every agent drives to exactly one customer ($\forall i : \sum_{j \in J} d(i, j) = 1$), and o_j agents drive to customer j (occupancy of j). An agent i can only carry a customer j , if i drives to j ($\forall i, j : c(i, j) \leq d(i, j)$). The customer j that yields the highest utility for some agent i is preferred by i (denoted as $p(i, j) = 1$). The number of agents preferring some customer j is denoted as p_j .

$$d(i, j) = \begin{cases} 1 & \text{if } i \text{ drives to } j, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{Definition 6.1})$$

$$\forall j : o_j = \sum_{i \in I} d(i, j) \quad (\text{Definition 6.2})$$

$$c(i, j) = \begin{cases} 1 & \text{if } i \text{ carries } j, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{Definition 6.3})$$

$$p(i, j) = \begin{cases} 1 & \text{if } \forall j' \in J \setminus \{j\} : u(i, j) \geq u(i, j') \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 6.4})$$

$$\forall j : p_j = \sum_{i \in I} p(i, j) \quad (\text{Definition 6.5})$$

The utility an agent i receives from a customer j $u(i, j)$ is determined as the weighted average of two components: The individual utility $u_i(i, j)$ represents the inverse distance between agent and customer. The shared utility $u_s(j)$ is the utility which is identical to all agents $i \in I$. $u_i(i, j)$ is a uniform distribution in the range between 0 and 1 independently calculated for every agent, $u_s(j)$ is a uniform distribution in the range between 0 and 1.

$$u(i, j) = \alpha \cdot u_i(i, j) + (1 - \alpha) \cdot u_s(j), 0 \leq \alpha \leq 1 \quad (\text{Definition 6.6})$$

The utilization fraction is calculated as the average number of agents carrying a customer (given by $f(i) = 1$) divided by the total number of agents N . The agent utilization $f(i)$ denotes if agent i carries any customer. The utility u is given by the average agent utility $u(i)$ which is 0 if agent i does not carry any customer and is $u(i, j)$ if i carries customer j .

$$f = \frac{1}{N} \cdot \sum_{i \in I} f(i) \quad (\text{Definition 6.7})$$

$$f(i) = \sum_{j \in J} c(i, j) \quad (\text{Definition 6.8})$$

$$u = \frac{1}{N} \cdot \sum_{i \in I} u(i) \quad (\text{Definition 6.9})$$

$$u(i) = \sum_{j \in J} u(i, j) \cdot c(i, j) \quad (\text{Definition 6.10})$$

For numerical experiments and simulations I assume that there are $N = 1000$ agents and customers. I further assume that $\alpha = 0.5$, resulting in the same influence for shared and individual utility. The individual utility is uniformly distributed between $\frac{1}{N}$ and $u_{max} = 1$. Every agent that is successful at the preferred customer receives on average u_m and every agent successful at a randomly chosen customer receives on average $u_{avg} = 0.5$. Without loss of generality, I further assume that customers are indexed

by their shared utility ($u_s(j) = \frac{j}{N}$). Though deterministic rather than random, this does not influence numerical results (the index j is no more than a theoretical construct which one can fit to the utilities). It simplifies calculations, as one can easily iterate through all customers with a higher (or lower) shared utility.

6.2. Theoretic Foundations

The maximum utility an agent can achieve may be lower than $u_{max} = 1$ as the utility is built as the weighted sum of two uniformly distributed variables with maximum u_{max} .

6.2.1. Probability of a Customer with a given Shared Utility yielding Maximum Utility

It is possible that there is no longer a single customer yielding maximum utility, but there can be multiple customers with the same utility. A customer is part of the set of top customers for some agent if there is no customer who returns a higher utility for this agent.

For simplicity, I first consider random integers for the individual component rather than a random permutation for the shared component (no duplicates). With this simplification, the probability that the utility retrieved from one customer is higher than the utility retrieved from another customer is independent of the utility yielded by all other customers (otherwise, one had to ensure that no duplicates occurred).

I denote the probability $\Pi(j)$ that some customer j with shared utility component $u_s(j)$ is among the customers with highest utility for any agent $i \in I$. Assuming that $u(i, j) = \alpha \cdot u_i(i, j) + (1 - \alpha) \cdot u_s(j)$ (Definition 6.6) and that $u_i(i, j)$ is random, I conclude that this probability only depends on the customer j .

$$\Pi(i, j) = \Pi(j) = P(\forall j' : u(i, j) \geq u(i, j')) \quad (29)$$

Without loss of generality, one can assume that $u_s(j) = \frac{j}{N}$. In the following I will use j as $u_s(j) \cdot N$. Numerically, I assume that every individual utility between $\frac{1}{N}$ and 1 is equally likely, I use $q \in 1 \dots N$ to model all possible individual utilities ($q = u_i(i, j) \cdot N$). I separately calculate the probability that another customer yields higher utility for those customers with a higher ($\Pi_h(j, q)$) and a lower ($\Pi_l(j, q)$) shared utility component. The total number of customers considered in $\Pi_l(j, q)$ and $\Pi_h(j, q)$ is $N - 1$, customers $j_l < j$ are considered in $\Pi_l(j, q)$, customers $j_h > j$ are considered in $\Pi_h(j, q)$.

$$\Pi(j) = \frac{1}{N} \sum_{q=1}^N \Pi_l(j, q) \Pi_h(j, q) \quad (30)$$

To derive the formulas for $\Pi_l(j, q)$ and $\Pi_h(j, q)$, I first consider a basic example: In an environment with $N = 5$ customers and agents, there is a customer $j = 3$ with shared utility $u_s(j) = \frac{3}{N}$ and an agent i assigning an individual utility $u_i(i, j) = \frac{3}{N}, q = 3$ to j . What is the probability that a customer with a lower shared utility $j_l \in \{1, 2\}$ or a higher shared utility $j_h \in \{4, 5\}$ is preferred over j by agent i ? Agent i can assign any individual utility $\frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, \frac{5}{5}$ to these customers $j' \in \{1, 2, 4, 5\}$ (resulting in $q' \in \{1, 2, 3, 4, 5\}$). For every customer j' one determines the probability that this customer does not reach a higher utility than $u(i, j) = \alpha \cdot u_i(i, j) + (1 - \alpha) \cdot u_s(j) = \frac{3}{N}$. In table 4 I display the (combined) utility of j' (multiplied by N for readability) and whether j or j' reaches a higher utility for agent i ($\rightarrow j$ and $\rightarrow j'$), depending on its individual utility q' that an agent i can derive from j' (leftmost column). The last row gives the probability that j' does not exceed j . As none of the other customers must reach a higher utility, I multiply the probabilities (that is $\frac{5}{5} \cdot \frac{4}{5} \cdot \frac{2}{5} \cdot \frac{1}{5} = \frac{8}{125}$) to retrieve the probability that customer j reaches the highest utility for agent i , if agent i assigned him an individual utility of $\frac{q}{N} = \frac{3}{5}$. Obviously, one has to calculate the probability that j is the highest utility customer for all possible individual utilities, that is all values of $q \in \{1 \dots N\}$.

$\Pi_l(j, q)$ is 1 if customer j has the lowest shared utility ($j = 1$) as there is no customer with a lower shared utility who could exceed the utility of customer j . Thus, j yields a higher utility than all customers with a lower shared utility. Otherwise, it is the product of the probabilities that the utility of j exceeds the utility of all customers $j' = j - j_l$ with a lower shared utility. The probability of exceeding any given other customer is given by $\frac{q+j_l}{N}$, but at most 1 ($\frac{N}{N}$). If a customer j' has a j_l lower shared utility than j , its individual utility must be at least $j_l + 1$ higher than the individual utility of j (q) to exceed j . I, therefore, calculate the probability that the individual utility of the other customer j' is not more than $q + j_l$.

$$\Pi_l(j, q) = \begin{cases} \prod_{j_l=1}^{j-1} \frac{\min(N, q+j_l)}{N}, & \text{if } j > 1 \\ 1 & \text{otherwise} \end{cases} \quad (31)$$

$\Pi_h(j, q)$ is 1 if customer j has the highest shared component as no customer with a higher shared utility component exceeds the utility of j . Otherwise, it is the product of the probabilities that the utility of j exceeds every customer $j' = j + j_h$ with a higher shared utility. The probability of exceeding a given other customer is given by $\frac{q-j_h}{N}$, but is always non-negative. If a customer j' has a shared utility that is j_h higher than the one of j , its individual utility must be at most $j_h - 1$ lower than the individual utility of j (q). j' , therefore, requires an indi-

Indiv. Utility $q' = u_i(i, j') \cdot N$	Lower		Higher	
	$j' = 1$	$j' = 2$	$j' = 4$	$j' = 5$
1	$1 \rightarrow j$	$1.5 \rightarrow j$	$2.5 \rightarrow j$	$3 \rightarrow j$
2	$1.5 \rightarrow j$	$2 \rightarrow j$	$3 \rightarrow j$	$3.5 \rightarrow j'$
3	$2 \rightarrow j$	$2.5 \rightarrow j$	$3.5 \rightarrow j'$	$4 \rightarrow j'$
4	$2.5 \rightarrow j$	$3 \rightarrow j$	$4 \rightarrow j'$	$4.5 \rightarrow j'$
5	$3 \rightarrow j$	$3.5 \rightarrow j'$	$4.5 \rightarrow j'$	$5 \rightarrow j'$
prob. $u(i, j) \geq u(i, j')$	$\frac{5}{5}$	$\frac{4}{5}$	$\frac{2}{5}$	$\frac{1}{5}$

Table 4: MP: Highest Utility Customer (Example)

vidual utility of $q - j_h + 1$ to exceed the utility of j . The combined utility is higher for j , if the individual utility of j' is at most $q - j_h$.

$$\Pi_h(j, q) = \begin{cases} \prod_{j_h=1}^{N-j} \frac{\max(0, q-j_h)}{N}, & \text{if } j < N \\ 1 & \text{otherwise} \end{cases} \quad (32)$$

Incorporating equations 31 and 32 in equation 30 yields:

$$\Pi(j) = \begin{cases} \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_l=1}^{j-1} \frac{\min(N, q+j_l)}{N}, & \text{if } 1 < j < N \\ \prod_{j_h=1}^{N-j} \frac{\max(0, q-j_h)}{N}, & \text{if } 1 < j < N \\ \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_h=1}^{N-1} \frac{\max(0, q-j_h)}{N}, & \text{if } j = 1 \wedge N \neq 1 \\ \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_l=1}^{N-1} \frac{\min(N, q+j_l)}{N}, & \text{if } j = N \wedge N \neq 1 \\ 1 & \text{otherwise} \end{cases} \quad (33)$$

This equation 33 can be transformed to the random permutation case by decreasing the denominator as the number of options for the individual component of the other customer is reduced by the assignment to the first customer. This also decreases the numerator of the fraction in Π_j .

$$\Pi(j) = \begin{cases} \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_l=1}^{j-1} \frac{\min(N-1, q+j_l-1)}{N-1}, & \text{if } 1 < j < N \\ \prod_{j_h=1}^{N-j} \frac{\max(0, q-j_h)}{N-1}, & \text{if } 1 < j < N \\ \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_h=1}^{N-1} \frac{\max(0, q-j_h)}{N-1}, & \text{if } j = 1 \wedge N \neq 1 \\ \frac{1}{N} \cdot \sum_{q=1}^N \prod_{j_l=1}^{N-1} \frac{\min(N-1, q+j_l-1)}{N-1}, & \text{if } j = N \wedge N \neq 1 \\ 1 & \text{otherwise} \end{cases} \quad (34)$$

Given this approach, it might happen that two customers yield the same utility. The probability that the highest utility is shared among different customers decreases with $N \rightarrow \infty$. For $N = 1000$, approximately 3.9% of all agents prefer more than one customer (given by the sum of probabilities $\Pi(j)$ for all j).

With the above equation with $N = 1000$, I expect that 4.03% of all agents prefer the customer with the highest shared utility (that is $\max(j)$). For those 70 customers with the highest shared component the probability of an agent preferring them is greater than 0.1%, thus, on average, there is an agent for whom this customer yields the best possible utility.

6.2.2. Expected Number of Agents Sharing a Top Priority

The number of agents sharing the same top priority customer depends on the shared component of this customer. The customer with the highest possible shared utility will be chosen more often than the customer with the lowest shared utility.

$$Pref(p_j) = \binom{N}{p_j} (\Pi(j))^{p_j} (1 - \Pi(j))^{N-p_j} \quad (35)$$

$$= \binom{N\Pi(j)}{p_j} \left(\frac{1}{N}\right)^{p_j} \left(1 - \frac{1}{N}\right)^{N\Pi(j)-p_j} \quad (36)$$

$$= \frac{(\Pi(j))^{p_j}}{p_j!} e^{-\Pi(j)} \quad (37)$$

6.2.3. Expected Number of Distinct Top Priorities

With the equation 37, it is now possible to calculate the probability that a customer is noone's first preference ($Pref(0)$ for $p_j = 0$) and the expected number of customers which are noone's preference (as the average probability).

I ignore duplicate first preferences and assume that a customer is selected with his associated probability of being first preference.

$$\text{no. of not pref. customers} = N - \sum_{j \in J} 1 - P_j(0) \quad (38)$$

The expected number of customers who are not preferred by any agent for $N = 1000$ is, therefore, 923 (or alternatively: I expect approximately 77 distinct first preferences).

6.2.4. Expected Utility of Top Priority Customers

The expected utility of a randomly selected customer is straight-forward: The average of two random numbers between $\frac{1}{N}$ and 1 is $u_{avg} = 0.5$ (for sufficiently large N). The expected utility for the first preference customer u_m is more elaborate: $u_m = u_{max} = 1$ can only be reached, if both the shared and the individual utility are maximum for an agent i and a customer j . Otherwise, the maximum agent utility is a weighted sum of $\frac{1}{N} \cdot \frac{j+q}{2}$ weighted by the probability that a customer yielding shared utility $\frac{j}{N}$ and individual utility $\frac{q}{N}$ (for agent i). For simplicity, I only consider the case $1 < j < N$; equation 39 needs to be adjusted accordingly to equation 34 to cater for $j = 1$ and $j = N$. For the defined numerical assumptions, the expected utility of top priority customers is $u_m = 0.92$.

$$u_m = \frac{1}{N} \sum_{q=1}^N \prod_{j_l=1}^{j-1} \frac{\min(N-1, q+j_l-1)}{N-1} \prod_{j_h=1}^{N-j} \frac{\max(0, q-j_h)}{N-1} \cdot \frac{j+q}{2} \quad (39)$$

6.3. No Learning

Agents incorporating the NL strategy randomly choose where to drive to. Thus, the number of agents per customer is Poisson distributed around 1. The number of agents carrying a customer equals the number of customers who are carried by some agent which is N minus the number of agents who are not carried by any agent ($\sum_{i \in I} c(i, j) = 0$). As the number of agents driving to some customer j is Poisson distributed, I conclude that the number of agents who do not carry any agent is $(1 - e^{-1}) \cdot N$, resulting in a utilization fraction of $f = f_{NL} = 63.2\%$ and a utility of $u = f \cdot u_{avg} = 0.316 \cdot u_{max}$.

6.4. Rank Dependent Choice

Obviously, only those customers who are some agent's first preference will be served with the RD strategy.

The utilization fraction is, therefore, given by equation 38 ($f = f_{RD} = 7.7\%$ for $N = 1000$). Those 7.7% of all agents will receive maximum utility, resulting in $u = f \cdot u_m = 0.075 \cdot u_{max}$ (with $u_m = 0.92$ from equation 39).

6.5. Limited Learning

In the LL strategy, agents decide randomly on a customer until they are able to serve one. After that, agents try their preferred customer. If they are being "bullied" out, they return to selecting randomly. Those customers who are preferred by some agent ($j \in J | \exists i \in I : p(i, j) = 1$) will be carried in all iterations unless they did not carry any customer in the previous iteration $t-1$ ($f_{t-1} \cdot f_{RD}$). Agents who do not drive to their preferred customer randomly select any customer, resulting in $(1 - e^{f_{t-1}-1})$ as the number of agents in this phase is lower than the number of customers to choose from.

$$f_t = f_{t-1} \cdot f_{RD} + (1 - e^{f_{t-1}-1}) \quad (40)$$

$$f = \lim_{t \rightarrow \infty} f_t \quad (41)$$

$$u = f \cdot f_{RD} \cdot u_m + (1 - e^{f-1}) \cdot (r \cdot u_m + (1-r) \cdot u_{avg}) \quad (42)$$

Derived from equation 41 and ?? (with f_{RD} the number of customers who are preferred by some agent (or the utilization fraction of the RD strategy), $u_m = 0.97$ and $r = \sum_{j \in J} \Pi(j) \cdot e^{-\Pi(j)} = 0.01$), I deduce that the utilization fraction is $f = 45.5\%$ and that the average utility is $u = 0.246 \cdot u_{max}$.

6.6. One Period Repetition

Using the OPR strategy, agents drive to their preferred customer after being successful with some randomly chosen customer for two iterations.

$f_{RD} \cdot N$ agents carry the customer they prefer (f_{RD} is the utilization fraction of the RD strategy). All other agents follow the three-step approach ((1) random, (2) return, and (3) improve). In every iteration a fraction $1 - 2x$ agents chooses randomly ($x = (1 - 2x) \cdot (1 - e^{-1})$ successful), x agents return, and x agents drive to their preferred customer (which is already occupied by another agent, therefore not increasing the utilization). For $N = 1000$ and, therefore, $f_{RD} = 0.077$, the utilization fraction is $f = 56.6\%$.

$$f = f_{RD} + (1 - f_{RD}) \cdot \left(x + (1 - 2x) \cdot (1 - e^{-1}) \right) \quad (43)$$

The utility is calculated analogously, those f_{RD} agents carrying their preferred customer receive $u_m = 0.97$, the other agents carry a randomly selected customer and, therefore, receive u_{avg} . This results in $u = 0.320 \cdot u_{max}$.

6.7. Crowd Avoiding

Agents who follow the CA strategy randomly choose any customer who was not carried during the previous iteration. Thus, there are N agents driving to $(1 - f_{t-1}) \cdot N$ customers.

$$\begin{aligned} f_t &= (1 - f_{t-1}) \cdot \left(1 - \binom{N}{0} \cdot \left(1 - \frac{1}{(1 - f_{t-1}) \cdot N} \right)^N \right) \\ &= (1 - f_{t-1}) \cdot \left(1 - e^{-(1-f_{t-1})} \right) \end{aligned} \quad (44)$$

I, therefore, conclude that $f = 45.7\%$. As all successful agents drive to a randomly chosen a customer, I assume that these agents receive u_{avg} . Thus, the utility is $u = 0.229 \cdot u_{max}$.

6.8. Stochastic Crowd Avoiding

Using the SCA strategy, agents either return to the same customer or drive to any other customer depending on the number of agents driving to the customer they drove to in the previous iteration. If at time $t - 1$ agent i drove to customer j ($d(i, j) = 1$) and the occupancy of customer j is $o_j = 1$, agent i returns to customer j at time t . If agent i drove to customer j at time $t - 1$ and the occupancy $o_j > 1$, i returns there with probability $\frac{1}{o_j}$ and randomly chooses any other customer at time t with probability $\frac{o_j - 1}{o_j}$.

In simulations with $N = 1000$, $u_{max} = 1$ and $u_{avg} = 0.5$ I observe a utilization fraction of $f = 73.5\%$ and a utility of $u = 0.368 \cdot u_{max}$.

6.9. Stochastic Rank Dependent Choice

With this strategy, the probability of driving to the top customer depends on the number of agents which share the same top priority.

Analytically, one can assume that the function of the utilization fraction has to incorporate the no longer random number of agents preferring some customer. $Pref(p_j)$ is the probability that a customer j is preferred by exactly p_j agents (derived from equation 37). $F(p_j)$ is the expected utilization, if p_j agents prefer customer j . As $Pref(p_j)$ is used to weight $F(p_j)$, one has to divide by $\sum_{j \in J} \sum_{p_j=1}^N Pref(p_j) = \sum_{j \in J} \Pi(j) \approx N$.

$$f = \frac{1}{\sum_{j \in J} \Pi(j)} \sum_{j \in J} \left(\sum_{p_j=1}^N Pref(p_j) \cdot F(p_j) \right) \quad (45)$$

$F(p_j)$ includes the probability that r_j agents divert to other customers (with probability $D(p_j, r_j) = ((p_j - r_j)!)^{-1} \cdot e^{-1}$).

$$F(p_j) = \sum_{r_j=1}^{p_j} D(p_j, r_j) \cdot s \cdot r_j' + \sum_{r_j''=0}^{p_j-1} D(p_j, r_j'') \quad (46)$$

The average utility is calculated by adapting equations 45 and 46 such that it incorporates different utility levels regarding on the agent's type of choice (remain with their top priority resulting in u_m or diverting to alternative resources resulting in u_{alt}).

$$u = \frac{1}{\sum_{j \in J} \Pi(j)} \sum_{j \in J} \left(\sum_{p_j=1}^N Pref(p_j) \cdot U(p_j) \right) \quad (47)$$

$$\begin{aligned} U(p_j) &= \sum_{r_j'=0}^{p_j-1} D(p_j, r_j') \cdot s \cdot r_j' \cdot u_{alt} + \\ &\quad \sum_{r_j''=1}^{p_j} D(p_j, r_j'') \cdot u_m \end{aligned} \quad (48)$$

s , u_m , and u_{alt} depend on the actual strategy. Table 5 compares the variables for SRD1 and SRD2.

6.9.1. Noone's First Choice Customer

In this strategy, agents choose those customers who are not preferred by any agent ($j \in J$, s.t. $\sum_{i \in I} p(i, j) = 0$).

As the number of diverting agents on average equals the number of customers who are not preferred by any agent, I can assume that a fraction of $s = 0.632$ of all diverting agents successfully carries another customer

Strategy	s	u_m	u_{alt}	f	u
SRD1	0.632	0.92	0.46	89.8%	0.521
SRD2	0.661	0.92	0.50	88.0%	0.512

Table 5: MP: SRD Strategy – Variables

(success rate). The utilization fraction is, therefore, $f = 63.8\%$. The utility of diverting agents (alternate utility) is u_{alt} . One cannot assume $u_{alt} = 0.5$, as only those customers with a lower shared component and therefore a lower utility are being selected as noone's preference. For $N = 1000$, I assume $u_{alt} = 0.46$, as on average 77 of the highest utility customers cannot be selected. The expected maximum utility is $u_m = 0.92$, according to equation 39, the utility is thus $u = 0.326 \cdot u_{max}$.

6.9.2. Any Other Customer

The SRD2 strategy dictates diverting agents to choose any other customer, regardless of the preferences of other agents or own preferences. The success rate $s = 0.611$ therefore derived from equation 1 with $\lambda = 1 - f_{RD}$ as $(1 - f_{RD})N$ agents divert to N customers. f_{RD} is the utilization fraction of the RD strategy and can be interpreted as the fraction of customers who can carry their preferred customer in the SRD strategy.

$$s = (1 - f_{RD}) \cdot \left(1 - e^{-\frac{1}{1-f_{RD}}}\right) \quad (49)$$

Thus, the expected utilization fraction is $f = 61.9\%$. All agents carrying their preferred customer ($i \in I, \text{s.t. } \forall j \in J : c(i, j) = p(i, j)$) can expect $u_m = 0.92$ (as in equation 39). Diverting agents can expect $u_{alt} = 0.462$. The expected average utility is $u = 0.330 \cdot u_{max}$.

6.9.3. Second Choice Customer

In the SRD3 strategy, diverting agents drive to the customer yielding them the second highest utility. For this strategy, the utilization rises only slightly in comparison to the RD strategy, as those 92.3% of all agents who randomly choose not to service the top ranked customer will go to the second ranked customer, which in most cases is someone else's top priority or overlaps with another agent's second priority.

The number of distinct second preferences is around 93 for $N = 1000$. Yet, many of these customers are some other agent's first preference. The expected number of customers which are either first or second preference is, therefore, ≈ 94 (in simulations).

Simulations suggest a utilization fraction of $\bar{f} = 9.4\%$ and an average utility of $u = 0.091 \cdot u_{max}$.

6.9.4. Best Vacant Customer

A similar explanation holds for the strategy SRD4 (Best Vacant Customer): Even if agents only turn to customers who are noone's first preference, they will most likely be competing there, as those customers will also be much alike.

The total number of distinct customers in the best vacant customer choice is approx. 74 with $N = 1000$. With ≈ 77 distinct first preference customers, there are around 151 customers the agents choose from.

The actual utilization is lower, as agents do not distribute themselves uniformly. In simulations, the utilization fraction was $\bar{f} = 12.1\%$ and the utility was $u = 0.115 \cdot u_{max}$.

6.10. Results

The utilization fraction and utility for all considered strategies can be found in table 6.

All strategies which do not incorporate the utility (NL, CA, SCA) are obviously not affected by mixed utilities. LL, OPR, RD, and SRD on the opposite worsen (moderately to dramatically) in comparison to the *Individual Preferences* setting. Only one of the rank dependent strategies outperforms both baseline comparisons: SRD1 (and with respect to utility OPR as well). As the redirection option for SRD3 and SRD4 is correlated to the first choice, and due to the low number of distinct first preferences, those strategies fall behind SRD1 and SRD2. With the decreased performance of rank dependent strategies (most "first preference selections" do not increase utility and utilization), SCA becomes the best strategy concerning both utilization fraction and utility.

7. Individual Preferences with Multiple Customers per District

In this model variant I assume that there are several customers in one district, thus, an agent always has several customers from which he can carry one even if the preferred one is not available. I assume that every district on average has the same number of customers, but as customers randomly spawn in some district, there can also be less or more customers in a district. Agents select a customer and drive to the district in which the selected customer is located in.

Strategy	utilization f	utility u
NL	63.2%	0.316
RD	7.7%	0.075
LL	45.5%	0.246
OPR	56.6%	0.320
CA	45.7%	0.229
SCA	73.5%	0.368
SRD1	63.8%	0.326
SRD2	61.9%	0.313
SRD3	9.4%	0.091
SRD4	12.1%	0.115

Table 6: MP: Comparing Strategies

7.1. The Model

In the IPMC model variant, customers are located to districts. Agents $i \in I, |I| = N$ drive to their preferred customer and are able to divert to other customers in the same district at no cost. I denote that some customer $j \in J, |J| = N$ is located in a district $k \in K, |K| = D = \frac{N}{\varphi}$ as $b(j, k) = 1$ (j "belongs to" k). Every customer j belongs to exactly one district k ($\forall j : \sum_{k \in K} b(j, k) = 1$), and c_k customers are located in district k (capacity of k).

$$b(j, k) = \begin{cases} 1 & \text{if } j \text{ is in } k \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 7.1})$$

$$\forall k : c_k = \sum_{j \in J} b(j, k) \quad (\text{Definition 7.2})$$

Agents drive to customers. I denote this relation as $d(i, j) = 1$. Every agent drives to exactly one customer ($\forall i : \sum_{j \in J} d(i, j) = 1$). As agents are able to divert to other customers in the same district, I extend $d(i, j) = 1$ as the notion that agent i drives to customer j to $d(i, k) = 1$ to denote that agent i drives to the district k that j is located in ($d(i, j) = 1 \wedge b(j, k) = 1 \Rightarrow d(i, k) = 1$). As the customer j that agent i originally drove to exactly one district k , I conduct that every agent drives to exactly one district ($\forall i : \sum_{k \in K} d(i, k) = 1$). The number of agents driving to some district k yields the occupancy o_k .

$$d(i, j) = \begin{cases} 1 & \text{if } i \text{ drives to } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 7.3})$$

$$\forall j : o_j = \sum_{i \in I} d(i, j) \quad (\text{Definition 7.4})$$

$$d(i, k) = \begin{cases} 1 & \text{if } i \text{ drives to } k \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 7.5})$$

$$\forall k : o_k = \sum_{i \in I} d(i, k) \quad (\text{Definition 7.6})$$

Agents can carry any customer that awaits a ride in the district k that agent i drove to. I denote that agent i carries customer j as $c(i, j) = 1$. One customer can only be carried by one agent ($\forall j : \sum_{i \in I} c(i, j) \leq 1$) and one agent i can carry at most one customer ($\forall i : \sum_{j \in J} c(i, j) \leq 1$). Agents can only carry customers located in the district they drove to ($c(i, j) \leq \sum_{k \in K} d(i, k) \cdot b(j, k)$). If agents are able to carry any customer, they prefer carrying him over not carrying anyone. Thus, the total number of customers carried from one district k is the minimum of the number of customers in k (capacity c_k) and the number of agents driving to k (occupancy o_k) ($\forall k : \sum_{i \in I} c(i, j) \cdot b(j, k) = \min(c_k, o_k)$).

$$c(i, j) = \begin{cases} 1 & \text{if } i \text{ carries } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 7.7})$$

Agents can either drive to their preferred customer or district or randomly choose a resource. I use $p(i, j) = 1$ to denote that i prefers j (j yields more utility for i than any other customer). This is the case if no other customer j' results in a higher utility. The number of agents preferring j is given as p_j . Analogously, I define p_k as the number of agents preferring any customer that are located in district k .

$$p(i, j) = \begin{cases} 1 & \text{if } \forall j' : u(i, j) \geq \\ & u(i, j') \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 7.8})$$

$$\forall j : p_j = \sum_{i \in I} p(i, j) \quad (\text{Definition 7.9})$$

$$\forall k : p_k = \sum_{i \in I} \sum_{j \in J} p(i, j) \cdot b(j, k) \quad (\text{Definition 7.10})$$

$u(i, j)$ is a random permutation individually assigned for every agent ($\forall i \in I : \forall j, j' \in J : u(i, j) = u(i, j') \Rightarrow j = j'$).

j'). As agents who select their preferred resource choose whichever customer results in the highest utility for them and drive to the corresponding district, I define that the utility of a district k is determined by the highest utility of any customer in k .

$$\forall k : u(i, k) = \max_{j \in J} (u(i, j) \cdot b(j, k)) \quad (\text{Definition 7.11})$$

One calculates the utilization fraction as the share of successful agents, that is agents who carry some customer. The utility is the average of all agent utilities $u(i)$. $u(i)$ is the utility agent i receives. If i does not carry any customer, the agent utility is $u(i) = 0$, otherwise it is the utility $u(i, j)$ of the customer j that agent i carries.

$$f = \frac{1}{N} \cdot \sum_{i \in I} f(i) \quad (\text{Definition 7.12})$$

$$f(i) = \sum_{j \in J} c(i, j) \quad (\text{Definition 7.13})$$

$$u = \frac{1}{N} \cdot \sum_{i \in I} u(i) \quad (\text{Definition 7.14})$$

$$u(i) = \sum_{j \in J} u(i, j) \cdot c(i, j) \quad (\text{Definition 7.15})$$

In simulations and numerical experiments, I assume that there are $N = 1000$ agents and customers in $D = 200$ districts (on average $\varphi = 5$ customers per district), that the utility is uniformly distributed between $\frac{1}{N}$ and $u_{max} = 1$. Every agent that is successful in the preferred district receives on average u_m and every agent successful at a randomly chosen district receives on average $u_{avg} = 0.5$.

7.2. Theoretic Foundations

7.2.1. Capacity: Number of Customers per District

In theory, there can be $0 \dots N$ customers in one district, though both extremes are highly unlikely. Assuming that there are φ customers on average per district ($N = \varphi D$), the probability $C(c_k)$ for capacity c_k is given by equation 50. In this case φ is the average number of customers per district (in numerical experiments and simulations: $\varphi = 5$).

$$\begin{aligned} C(c_k) &= \binom{\varphi D}{c_k} \cdot \left(\frac{1}{D}\right)^{c_k} \cdot \left(1 - \frac{1}{D}\right)^{\varphi D - c_k} \\ &= \frac{\varphi^{c_k}}{c_k!} \cdot e^{-\varphi} \end{aligned} \quad (50)$$

7.2.2. Occupancy: Number of Agents per District (based upon Capacity)

As agents choose a customer and then drive to the corresponding district, the probability that o_k agents drive to district k depends on its capacity c_k . With N agents and N customers, the number of agents in district k with c_k customers is Gaussian distributed around c_k .

$$O(o_k, c_k) = \frac{c_k^{o_k}}{o_k!} e^{-c_k} \quad (51)$$

7.2.3. Same First Preference

The probability that a district with capacity c_k is preferred by p_k agents is calculated as a Gaussian distribution around c_k , as agents randomly "choose" their preferred customer.

$$Pref(c_k, p_k) = \frac{c_k^{p_k}}{p_k!} \cdot e^{-c_k} \quad (52)$$

7.2.4. Expected Utility of Top Priority Customers

The expected maximum utility depends on the capacity c_k : If an agent i enters a district with c_k customers and he carries any customer in this district, there is a $\frac{1}{c_k}$ chance that the customer j that i carries is his preferred customer yielding a utility of u_{max} and a $1 - \frac{1}{c_k}$ chance that i carries any other customer, yielding a utility of on average u_{avg} .

$$u_m(c_k) = \frac{1}{c_k} \cdot u_{max} + \frac{c_k - 1}{c_k} \cdot u_{avg} \quad (53)$$

The expected maximum utility u_m in random processes is calculated by weighting $u_m(c_k)$ by the probability of c_k and the expected number of successful agents $o_k \leq c_k$. I, therefore, conclude $u_m = 0.59$ if c_k of district k is unknown.

7.3. No Learning

In this strategy, every agent randomly decides which district he will go to by randomly selecting a customer j and driving to the district k that j is located in. The agent is then randomly assigned a customer from the selected district. If there are less or equal agents than customers ($o_k \leq c_k$), every agent will be assigned a customers. Otherwise, there is a $\frac{c_k}{o_k}$ probability for every agent to actually be assigned a customer.

Agents select customers and drive to the corresponding districts rather than districts directly, as this increases the utilization fraction and utility, as every district is – on average – chosen by as many drivers as it can cater (instead of φ drivers on average per district). In the appendix I calculate the utilization fraction and utility for

district-based choice (??). To derive the utilization fraction, I calculate the expected number of not carried customers for every possible capacity c_k ($\sum_{o_k=0}^{c_k-1} O(o_k, c_k) (c_k - o_k)$) and derive the number of carried customers from it. The probability of capacity c_k is derived from equation 50 and the probability of occupancy o_k is derived from equation 51.

$$f = \frac{1}{\varphi} \sum_{c_k=1}^N C(c_k) \cdot \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, c_k) (c_k - o_k) \right) \tag{54}$$

$$= \frac{1}{\varphi} \sum_{c_k=1}^N \frac{\varphi^{c_k}}{c_k!} e^{-\varphi} \cdot \left(c_k - \sum_{o_k=0}^{c_k-1} \frac{c_k^{o_k}}{o_k!} e^{-c_k} (c_k - o_k) \right) \tag{55}$$

The utilization fraction is, therefore, $f = f_{NL} = 83.0\%$. With average utility for all successful agents, the expected utility is $u = f \cdot u_{avg} = 41.5\%$ for $N = 1000$.

7.4. Rank Dependent Choice

I now consider the strategy in which every agent drives to the district which provides him with the best possible utility that is the district containing the customer yielding the highest utility. There are different possible approaches to choosing the best district: Choose the district with the highest average utility from all customers in this district or choose the district which contains ones (individual) #1 priority customer. The first corresponds to selecting a district in *No Learning*, the second to selecting a customer. I only consider the latter as it results in a higher utilization and utility. Yet, one can find some insight on the first in the appendix.

The utilization fraction is the same as for the NL strategy (given by equation 55), as the preferred customer is randomly selected (resulting in $f = f_{RD} = 83.0\%$). The utility increases slightly in comparison to *No Learning*, as the probability of serving the top priority customer is increased.

$$u = \frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \left(c_k - \sum_{o_k=0}^{c_k} O(o_k, c_k) (c_k - o_k) \right) \cdot u_m(c_k) \tag{56}$$

For $N = 1000$ and $\varphi = 5$ this results in an average utility of $u = 0.495 \cdot u_{max}$.

7.5. Limited Learning

In the LL strategy, every agent first chooses a customer at random and – after carrying a customer – continues with the highest ranked district. With multiple customers

in a district, one has to choose which district one deems #1 priority (district containing highest utility customer).

The utilization fraction f depends on the fraction of agents servicing their top district for the first time and the fraction of agents who either randomly choose a district or return to the best possible district. From equation 55 I derive $f_{RD} = 83.0\%$ which is the fraction of customers carried by an agent preferring them, f_t is calculated iteratively. On average $(1 - f_{t-1}) \cdot N$ customers are not carried by first agents choosing their preferred customer the first time (and thus belong to the first summand of the equation). Thus, on average $\lambda = (1 - f_{t-1})$ customers per district are not carried by agents belonging to the left summand of the equation.

$$f_t = f_{t-1} \cdot f_{RD} +$$

$$\underbrace{\frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda) \cdot (c_k - o_k) \right)}_{\text{random or return}} \tag{57}$$

f converges towards $f = 85.2\%$ for $f_{RD} = 0.830$.

To calculate the utility u , I adapt equation 57 to incorporate whether agents expect maximum utility $u_m(c_k)$ or average utility u_{avg} . All those agents who carry a customer from their highest utility district receive on average $u_m(c_k)$. As the right half of the equation comprises both those agents who randomly choose any resource and those, who return to their highest utility customer, I have to differentiate between those groups by introducing r as the fraction of agents returning to their highest utility resource. r is calculated as the fraction of customers in not overutilized districts ($r = \sum_{c_k=1}^N \sum_{p_k=0}^{c_k} Pref(p_k, c_k) = 0.621$). Thus, I derive $u = 0.500 \cdot u_{max}$.

$$u = f \cdot f_{RD} \cdot u_m(c_k) + \frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda) \cdot (c_k - o_k) \right) \cdot (r \cdot u_m(c_k) + (1 - r) \cdot u_{avg}) \tag{58}$$

7.6. One Period Repetition

Agents applying the OPR strategy choose the district containing their top priority customer after returning once to a successful random district choice.

Drawing upon the results from section 5.6 I calculate the utilization fraction and the utility as follows. f_{RD}

agents carry a customer from their preferred district, all other agents follow a three step approach: (1) random choice (with a success probability of f_{NL}), (2) return to the same district (certainly successful, that is $f(i) = 1$, as randomly choosing agents only drive to previously not carried customers), and (3) try best district (with a success rate of 0, as the agent would otherwise belong to those f_{RD} agents who are constantly successful). In every iteration, a share x of all agents is in step (2) and (3), and a share of $1 - 2x$ is in step (1) (successful with probability f_{NL} , resulting in $x = (1 - 2x) \cdot f_{NL} \approx 0.312$). f_{RD} is the utilization of the RD strategy and f_{NL} is the utilization of the NL strategy.

$$f = (x + (1 - 2x) \cdot f_{NL}) \cdot (1 - f_{RD}) + f_{RD} \quad (59)$$

$$u = (x + (1 - 2x) \cdot f_{NL}) \cdot (1 - f_{RD}) \cdot u_{avg} + u_m \cdot f_{RD} \quad (60)$$

Thus, I expect a utilization fraction of $f = 93.6\%$. The average utility is $u = 0.547 \cdot u_{max}$.

7.7. Crowd Avoiding

Using the strategy CA, agents only choose from customers which have not been carried the previous time step and drive to the district the selected customer is located in. This yields a weighted selection of the districts with too few agents. The number of customers which can be chosen at some time t is the number of customers not chosen at time $t - 1$. Those remaining customers are located in different districts. On average, a fraction of $\lambda = \frac{1}{1-f}$ of all customers remain vacant. I assume that these remaining customers are Gaussian distributed across districts, resulting in $\lambda \cdot c_k$ customers remaining per district.

$$f = (1 - f) \cdot \left(\sum_{c_k=1}^N C(c_k) \cdot \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda \cdot c_k) \cdot (c_k - o_k) \right) \right) \quad (61)$$

With the above assumptions, one can derive $f = 49.7\%$. As all agents randomly decide upon a resource, I conduct $u = 0.249 \cdot u_{max}$.

7.8. Stochastic Crowd Avoiding

With this strategy, agents deterministically return to the same district, if the capacity of district was not exceeded in the previous iteration. Otherwise, agents stochastically return to the same district or drive to any other district.

There are two different choice mechanisms: Returning if the customer is not taken by others or returning if

the district has remaining capacity. In the appendix, I introduce a customer-based decision but will continue with a district-based decision in this chapter.

If the number of agents in a district does not exceed the number of customers, this agent will return there. Otherwise, the agent will move towards another customer with $p = 1 - \frac{c_k}{o_k}$ and return to the same district with $p = \frac{c_k}{o_k}$. The customer is then chosen at random from all available customers. In simulations, the utilization fraction is $\bar{f} = 93.8\%$. The utility is average for all agents serving a customer that time step and, therefore, $u = 0.469 \cdot u_{max}$ for $N = 1000$.

7.9. Stochastic Rank Dependent Choice

This strategy vastly builds upon the strategy *Rank Dependent Choice*. Yet, all those drivers who prefer an overcrowded district will not carry a customer with a given probability. With *Stochastic Rank Dependent Choice*, these drivers are now diverted to another district with some probability $p = \frac{c_k - p_k}{p_k}$. The district to divert to is either a district which has remaining capacity, any other district, the #2 district, or the highest utility district which has remaining capacity. The overall utilization fraction f is calculated as a generalization of equation 18.

$$f = \sum_{c_k=0}^N C(c_k) \cdot \sum_{p_k=0}^N Pref(p_k, c_k) \cdot F(c_k, p_k) \quad (62)$$

$$u = \sum_{c_k=0}^N C(c_k) \cdot \sum_{p_k=0}^N Pref(p_k, c_k) \cdot U(c_k, p_k) \quad (63)$$

$Pref(p_k, c_k)$ is the probability that p_k agents prefer a district with capacity c_k (equation 52). $C(c_k)$ is the probability that the capacity of some district k is c_k (given by equation 50). The utilization fraction function $F(c_k, p_k)$ calculates the expected utilization, if p_k agents prefer a district k with capacity c_k (including r_k agents redirecting to other districts with probability $D(c_k, p_k, r_k)$).

$$F(c_k, p_k) = \begin{cases} p_k & \text{if } p_k \leq c_k \\ \sum_{r_k=0}^{c_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k + c_k) + \sum_{r_k=c_k+1}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k + (p_k - r_k)) & \text{otherwise} \end{cases} \quad (64)$$

$$\begin{aligned}
 D(c_k, p_k, r_k) &= \binom{p_k}{r_k} \cdot \left(\frac{p_k - c_k}{p_k}\right)^{r_k} \cdot \left(1 - \frac{p_k - c_k}{p_k}\right)^{p_k - r_k} \\
 &= \frac{(p_k - c_k)^{r_k}}{r_k!} \cdot e^{c_k - p_k} \tag{65}
 \end{aligned}$$

The success rate s depends on the strategy and its associated behavior in case of swapping.

The utility function $U(c_k, p_k)$ is given by adapting equation 64 accordingly to equation 22:

$$U(c_k, p_k) = \begin{cases} p_k \cdot u_m & \text{if } p_k \leq c_k \\ \sum_{r_k=0}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k \cdot u_{alt} + c_k \cdot u_m) & \\ + \sum_{r_k=0}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k \cdot u_{alt} + (p_k - r_k) \cdot u_m) & \text{otherwise} \end{cases} \tag{66}$$

Table 7 lists the variables s , u_{max} , and u_{alt} for the different SRD strategies.

In strategies SRD1 and SRD4, I assume that $s = 0.595$ as given by equation 67. On average $0.17N = (1 - 0.83)N$ agents divert to other districts. Thus, $0.17N$ customers are not being serviced by an agent to whom they are first preference. I furthermore assume that these customers are Gaussian distributed across all districts.

$$s = \sum_{c_k=1}^N \frac{\varphi^{c_k}}{c_k!} \cdot e^{-\varphi} \sum_{o_k=0}^{c_k-1} \frac{c_k^{o_k}}{o_k!} \cdot e^{-c_k}, \varphi = 5 \cdot 0.17 \tag{67}$$

In strategies SRD2 and SRD3, the success rate is $s = 0.442$. In this case, I calculate the expected number of previously not serviced customers ($c'_k = c_k - o_k + r_k$) and the probability that these customers are serviced by r'_k agents who divert to district k .

$$\begin{aligned}
 s &= \sum_{c_k=1}^N \sum_{c'_k}^{c_k} P(c'_k) \\
 &\left(c'_k - \sum_{r'_k=0}^{c'_k-1} \frac{((1 - f_{RD}) \cdot c_k)^{r'_k}}{r'_k!} \cdot e^{-((1 - f_{RD}) \cdot c_k)} \right) \tag{68}
 \end{aligned}$$

The utility u_m is derived from section 7.2.4. In strategies SRD3 and SRD4 I also use this value u_m for u_{alt} (the alternative choice utility), for SRD1 and SRD2 I set $u_{alt} = u_{avg}$.

7.10. Results

Table 8 lists utilization and utility for all disused strategies for the IPMC model variant.

In the IPMC setting, OPR outperforms all other strategies regarding the utility and is outperformed by SCA concerning f by only a slight margin. All strategies except CA exceed the utilization of the baseline comparisons NL and RD, with respect to utility, SCA also falls behind RD (and RD outperforms NL). I assume that a higher average number of customers per district φ further increases the numbers for utilization and utility, this comparison is, therefore, purely relative. In comparison to the previously presented IP and MP model variants, the utility values for different strategies in the IPMC models are close to each other, as average utility and expected utility of a top priority customer are rather close.

8. Mixed Preferences with Multiple Customers per District

8.1. The Model

In the MPMC model, customers are located in districts (“belong to”) and the utility consists of a customer-specific (“shared”) component and an “individual” component that is based on customer and agent. The shared utility models the payoff an agent receives from carrying a customer. All agents would receive the same payoff if they carried this customer. The individual component models the costs to get to the pickup location which is identical for all customers in one district but varies between different agents.

In the MPMC model, customers $j \in J, |J| = N$ are “clustered” in districts $k \in K, |K| = D = \frac{N}{\varphi}$. One average φ customers await a driver in one district. As customers are located in a randomly drawn district, the number of customers in a district is Gaussian-distributed around φ . Customers $j \in J$ belong to the district $k \in K$ in which they await a driver. Let’s denote this as $b(j, k) = 1$. Every agent is located in exactly one district ($\forall j : \sum_{k \in K} b(j, k) = 1$) and the number of customers that are located in a district k is its capacity c_k .

$$b(j, k) = \begin{cases} 1 & \text{if } j \text{ is in } k \\ 0 & \text{otherwise} \end{cases} \tag{Definition 8.1}$$

$$\forall k : c_k = \sum_{j \in J} b(j, k) \tag{Definition 8.2}$$

Agents $i \in I, |I| = N$ select customers $j \in J$ ($d(i, j) = 1$) and drive to the district k that j is located in. Every agent drives to exactly one customer ($\forall i : \sum_{j \in J} d(i, j) = 1$), and the number of agents driving to customer j is denoted as occupancy o_j . In the MPMC model, agents can divert

Strategy	s	u_m	u_{alt}	f	u
SRD1	0.595	0.59	0.50	89.8%	0.521
SRD2	0.442	0.59	0.50	88.0%	0.512
SRD3	0.442	0.59	0.59	88.0%	0.519
SRD4	0.595	0.59	0.59	89.8%	0.530

Table 7: IPMC: SRD Strategy – Variables

Strategy	utilization f	utility u
NL	83.0%	0.415
RD	83.0%	0.495
LL	85.2%	0.500
OPR	93.6%	0.547
CA	49.7%	0.249
SCA	93.8%	0.469
SRD1	89.8%	0.521
SRD2	87.2%	0.508
SRD3	87.2%	0.515
SRD4	89.8%	0.530

Table 8: IPMC: Comparing Strategies

to other customers that belong to the same district at no cost; I, therefore, extend $d(i, j)$ to $d(i, k)$ to denote that agent i drives to district k . Every agent i drives to exactly one district k ($\forall i : \sum_{k \in K} d(i, k) = 1$). If an agent drives to a customer j , he also drives to the district k that j belongs to ($d(i, j) = 1 \wedge b(j, k) = 1 \Rightarrow d(i, k)$). The occupancy o_k of district k is the number of agents i driving to k .

$$d(i, j) = \begin{cases} 1 & \text{if } i \text{ drives to } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 8.3})$$

$$\forall j : o_j = \sum_{i \in I} d(i, j) \quad (\text{Definition 8.4})$$

$$d(i, k) = \begin{cases} 1 & \text{if } i \text{ drives to } k \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 8.5})$$

$$\forall k : o_k = \sum_{i \in I} d(i, j) \quad (\text{Definition 8.6})$$

As agents independently decide upon the customer or district they drive to, distributions in which too many agents drive to some customers and too few customers drive to some other agents can and do frequently occur. I further introduce the notion $c(i, j) = 1$ to denote that agent i carries customer j . An agent i can carry a customer j , if i drives to the district k that j belongs to ($c(i, j) \leq \sum_{k \in K} d(i, k) \cdot b(j, k)$). One agent i can carry at most one customer j ($\forall j : \sum_{i \in I} c(i, j) \leq 1$) and one customer j can be carried by at most one agent i

($\forall i : \sum_{j \in J} c(i, j) \leq 1$). In every district, agents carry as many customers as possible, no agent refuses to carry a customer remaining at this district. Thus, the number of customers carried per district is either capacity c_k or occupancy o_k ($\forall k : \sum_{i \in I} c(i, j) \cdot b(j, k) = \min(c_k, o_k)$).

$$c(i, j) = \begin{cases} 1 & \text{if } i \text{ carries } j \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 8.7})$$

Agents can either drive to their preferred customer or a randomly drawn customer (given by the strategy). For every agent i there exists a customer j whom he prefers over all other customers, as this customer yields the highest utility for him. A customer j is preferred by p_j agents. Agents prefer the district their preferred customer belongs to. A district k is preferred by p_k agents.

$$p(i, j) = \begin{cases} 1 & \text{if } \forall j' : u(i, j) \geq u(i, j') \\ 0 & \text{otherwise} \end{cases} \quad (\text{Definition 8.8})$$

$$\forall j : p_j = \sum_{i \in I} p(i, j) \quad (\text{Definition 8.9})$$

$$\forall k : p_k = \sum_{i \in I} \sum_{j \in J} p(i, j) \cdot b(j, k) \quad (\text{Definition 8.10})$$

The utility an agent i can gain from carrying customer j depends on both an individual and a shared utility

component $(u_i(i, j), u_s(j) = u_s(i, j) \forall i)$. Both utilities are uniformly distributed between 0 and 1.

$$u(i, j) = \alpha \cdot u_i(i, j) + (1 - \alpha) \cdot u_s(j), 0 \leq \alpha \leq 1$$

(Definition 8.11)

$$\forall j, j' \in J : \forall k \in K : b(j, k) = b(j', k) \Rightarrow u_s(j) = u_s(j')$$

(Definition 8.12)

In the MPMC game model, the individual utility is identical for all customers which are located in a given district as the driving distance between agent and customer is identical for all customers in the same location (district).

$$\forall k \in K : u_i(i, j) = u_i(i, k) \vee b(j, k) = 0$$

(Definition 8.13)

I define that the utility of a district k is given by the utility of the customer yielding the highest utility (see Proposition 8.2.2). The highest utility customer is defined as $b_1(j, k) = 1$. Obviously, the "best" customer j (customer with highest utility) must be located in district k , and there must not be any other customer j' that also belongs to k that yields a higher shared utility.

$$\forall k : u(i, k) = \max_{j \in J} (u(i, j) \cdot b(j, k))$$

(Definition 8.14)

$$b_1(j, k) = \begin{cases} 1, & \text{if } b(j, k) = 1 \wedge \\ & (u_s(j) \geq u_s(j') \vee b(j', k) = 0 \forall j') \\ 0, & \text{otherwise} \end{cases}$$

(Definition 8.15)

The utilization fraction is calculated as the average of all agent utilizations. The agent utilization $f(i)$ defines whether an agent i carries any customer. The utility is calculated as the average of all agent utilities $u(i)$. $u(i)$ is 0, if i does not carry any customer and the utility of the customer j that i carries ($u(i, j)$) otherwise.

$$f = \frac{1}{N} \cdot \sum_{i \in I} f(i)$$

(Definition 8.16)

$$f(i) = \sum_{j \in J} c(i, j)$$

(Definition 8.17)

$$u = \frac{1}{N} \cdot \sum_{i \in I} u(i)$$

(Definition 8.18)

$$u(i) = \sum_{j \in J} u(i, j) \cdot c(i, j)$$

(Definition 8.19)

For numerical experiments and simulations I assume that there are $N = 1000$ agents and customers in $D = 200$ districts (on average $\varphi = 5$ customers per district), that $\alpha = 0.5$, that the individual utility is uniformly distributed between $\frac{1}{N}$ and $u_{max} = 1$ (with step size $\frac{1}{D}$, as the individual utility is calculated on a district basis) and every agent that is successful at the preferred customer receives on average u_m and every agent successful at a randomly chosen customer receives on average $u_{avg} = 0.5$, and that customers are indexed by their utility ($u_s(j) = \frac{j}{N}$).

8.2. Theoretic Foundations

8.2.1. Capacity: Number of Customers per District

The capacity c_k that is the number of customers belonging to district k is given as a Gaussian distribution around the average number of customers per district φ , as customers randomly choose the district they belong to. Thus, the probability for capacity c_k is calculated as follows:

$$C(c_k) = \binom{\varphi D}{c_k} \cdot \left(\frac{1}{D}\right)^{c_k} \cdot \left(1 - \frac{1}{D}\right)^{\varphi D - c_k}$$

$$= \frac{\varphi^{c_k}}{c_k!} \cdot e^{-\varphi}$$

(69)

8.2.2. Highest Utility Customer and District

Proposition: In the MPMC partial game model, agents only prefer the customer j with the highest shared utility in district k . If another customer j' who belongs to the same district k has a higher shared utility, j is not preferred by any agent.

Proof. Assume that $j, j' \in J$ are customers, $k \in K$ is the district both customers belong to such that $b(j, k) = 1$ and $b(j', k) = 1$. Assume that j a higher utility than j' ($u(i, j) < u(i, j')$). An agent i chooses the district which yields the highest utility, assume that this district is k ($p(i, k) = 1$). Thus, $\forall k' \in K \setminus \{k\} : u(i, k) \geq u(i, k')$. From definition Definition 8.14 I know that the utility of a district is given by the highest utility of any of the customers belonging to it. I assume that this customer is j .

$$u(i, j) > u(i, j')$$

| with Definition 8.11 (70)

$$\alpha \cdot u_i(i, j) + (1 - \alpha) \cdot u_s(i, j) >$$

$$\alpha \cdot u_i(i, j') + (1 - \alpha) \cdot u_s(i, j')$$

| with Definition 8.13 (71)

$$\alpha \cdot u_i(i, k) + (1 - \alpha) \cdot u_s(i, j') > \alpha \cdot u_i(i, k) + (1 - \alpha) \cdot u_s(i, j') \mid - \alpha \cdot u_i(i, k) \tag{72}$$

$$u_s(i, j) > u_s(i, j') \tag{73}$$

The probability that a customer j yields the highest utility in his district k (is the “best” customer) is denoted as $B_1(j, c_k)$ and is calculated as the probability that all customers j_h with a higher shared utility $u_s(j_h) > u_s(j)$ choose other districts ($\forall j_h : b(j_h, k) = 0$), j belongs to k ($b(j, k) = 1$) and exactly $c_k - 1$ customers j_l with lower shared utility choose this district k . Without loss of generality, I assume that there are $N - j$ customers with higher shared utility and $j - 1$ customers with lower shared utility (one assigns the identifiers j to customers based on their shared utility component). N is the number of customers and the number of agents ($|I| = |J| = N$), and D is the number of districts ($|K| = D = \frac{N}{\varphi}$).

$$B_1(j, c_k) = \underbrace{\binom{j}{c_k} \frac{1}{D} \frac{c_k D - 1}{D} \frac{D - 1^{j-c_k}}{D}}_{j_l \leq j} \underbrace{\frac{D - 1^{N-j}}{D}}_{j_h > j} = \binom{j}{c_k - 1} \frac{1}{D} \frac{c_k D - 1}{D} \frac{D - 1^{N-j}}{D} = \frac{\left(\frac{j}{D}\right)^{c_k}}{(c_k)!} \cdot e^{-\frac{j}{D}} \tag{74}$$

If the capacity of (another) district is unknown, one can use a generalization of equation 74. $B_1(j)$ ensures that all customers j_h with a higher shared utility component choose other districts and all those j_l with lower shared utility component are being ignored.

$$B_1(j) = \binom{N-j}{0} \frac{1}{D} \left(1 - \frac{1}{D}\right)^{N-j} = \frac{D - 1^{N-j}}{D} = \frac{N - 1^{N-j}}{\frac{N}{\varphi}} \tag{75}$$

8.2.3. Same First Preference

In the MPMC model, the probability that district k yields maximum utility is no longer equal for all $k \in K$, as the utility depends on a shared component all agents agree upon.

The average number of agents choosing a district k with individual utility $u_i(k) = u_i(j)$ is denoted as $\varphi \cdot \Pi'(k)$. $\Pi'(k)$ is calculated as the product of probabilities that no other customer j_l, j_h yields a higher utility $u(i, j_l), u(i, j_h)$ for any agent i and is best in his district for all customers $j \in J$.

$$\Pi'(j) = \prod_{j_l=1}^{j-1} P\left(u(i, j) \geq u(i, j_l) \vee \sum_{k=1}^D b_1(j_l, k) = 0\right) \cdot \prod_{j_h=j+1}^N P\left(u(i, j) \geq u(i, j_h) \vee \sum_{k=1}^D b_1(j_h, k) = 0\right) \tag{76}$$

Numerically, I adapt equation 76 as follows: I iterate through all customers with lower $j' = j - j_l$ and higher $j' = j + j_h$ shared utility component assuming $\forall j \in J : u_s(j) = \varphi \frac{k}{D} \vee b(j, k) = 0$, and weighting individual and shared utility component equally ($\alpha = 0.5$). A customer $j' = j - j_l$ (shared utility is $j_l \cdot \frac{1}{N}$ lower if an agent carries j' than if he carried j) does not exceed the utility of j if its individual utility is less than $(j_l - 1) \cdot \frac{1}{\varphi}$ higher. Assuming that individual utilities are represented by q ($u_i(i, k) = \varphi \cdot q \cdot \frac{1}{N}$), one can derive that the individual utility of the district k' that j' is located in must not be higher than $\varphi q + j_l$. Analogously, the individual utility of a customer j exceeds the utility of $j' = j + j_h$ (customer with higher shared utility) if the individual utility is correspondingly lower that is lower by $\varphi \cdot q - j_h$. If j' does not yield the highest shared utility in its district, I do not consider it.

$$\Pi'_l(j, \varphi, q) = \prod_{j_l=1}^{j-1} \left(1 - \underbrace{(B_1(j - j_l))}_{j' \text{ best}} \cdot \underbrace{\left(1 - \frac{\min(N, \varphi \cdot q + j_l)}{N} \right)}_{\substack{u(i, j) \geq u(i, j') \\ u(i, j) < u(i, j')}} \right)$$

$j' \text{ not best or } u(i, j) \geq u(i, j')$

$$\begin{aligned} \Pi'_h(j, \varphi, q) &= \prod_{j_h=1}^{N-j} (1 - B_1(j + j_h)) \\ &\cdot \left(1 - \frac{\max(0, \varphi \cdot q - j_h)}{N} \right) \\ \Pi''(j) &= \sum_{q=1}^D \Pi'_l(j, \varphi, q) \cdot \Pi'_h(j, \varphi, q) \\ \Pi'(j) &= \frac{D}{\sum_{j''=1}^N \Pi''(j'')} \cdot \Pi''(j) \end{aligned} \tag{77}$$

I, therefore, expect $\varphi \Pi'(j) = \varphi \Pi'(k)$ ($b_1(j, k) = 1$) agents preferring the district k in which j is the highest utility customer. Yet, the actual number of agents preferring k is Gaussian distributed around $\varphi \Pi'(k)$. $Pref(p_k, \varphi \Pi'(j))$ is the probability that district k with

the highest utility customer j is preferred by exactly p_k agents.

$$Pref(p_k, \varphi \Pi'(j)) = \frac{(\varphi \Pi'(j))^{p_k}}{p_k} \cdot e^{-\varphi \Pi'(j)} \quad (78)$$

8.2.4. Occupancy

The occupancy of district k depends on the type of choice: If agents decide randomly, the average number of agents in district k is its capacity c_k , otherwise, it is the expected number of agents preferring it ($\varphi \Pi'(k)$). In the following, λ_k is the expected number of agents driving to district k .

$$O(o_k, \lambda_k) = \frac{\lambda_k^{o_k}}{o_k} \cdot e^{-\lambda_k} \quad (79)$$

8.2.5. Expected Utility of Top Priority Customers

In the MPMC setting, all agents agree upon the same "best" customer inside a district (Proposition 8.2.2).

To calculate the expected agent utility, I assume that every customer in a district with capacity c_k and highest utility customer j yields on average $u_e(j, k)$ to the agent carrying him. $\bar{u}_i(j)$ is the average individual utility of the district k that j is located in.

$$u_e(j, k) = \frac{1}{c_k} (u_s(j) + \bar{u}_i(j)) + \frac{c_k - 1}{c_k} \left(\frac{u_s(j)}{2} + \bar{u}_i(j) \right) \quad (80)$$

$$\bar{u}_i(j) = \frac{1}{\Pi'(j)} \cdot \sum_{k=1}^D \underbrace{\varphi \cdot k}_{\text{utility}} \cdot \underbrace{\Pi'_l(j, \varphi, k) \cdot \Pi'_h(j, \varphi, k)}_{\text{if successful}} \quad (81)$$

The average utility of an agent who carries a customer from his preferred district $u_m = 0.785$ is a weighted average of all possible $u_e(j, k)$ (weighted by the probability $B_1(j, c_k)$).

8.3. No Learning

Using the NL strategy, all agents drive to a randomly selected customer. Thus, individual utility levels are irrelevant. The utilization fraction depends on (1) the capacity c_k of district k (associated with probability $C(c_k)$) and (2) the occupancy o_k of district k (associated with probability $O(o_k, c_k)$). $C(c_k) = \frac{\varphi^{c_k}}{c_k!} \cdot e^{-\varphi}$ is the probability that c_k customers are randomly assigned the same district given an average of φ customers per district. $O(o_k, c_k) = \frac{c_k^{o_k}}{o_k!} \cdot e^{-c_k}$ is the probability of o_k agents randomly driving to district k containing c_k customers.

$$f = \frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \cdot \left(c_k - \underbrace{\sum_{o_k=0}^{c_k-1} O(o_k, c_k) \cdot (c_k - o_k)}_{\text{expected remaining capacity}} \right) \quad (82)$$

Thus, the utilization fraction is $f = f_{NL} = 83.0\%$. As all agents decide randomly where to drive to, all successful agents will receive average utility u_{avg} . The utility is, therefore, $u = 0.415 \cdot u_{max}$.

8.4. Rank Dependent Choice

In the MPMC model, for every customer j , I calculate the average number of agents driving there if this customer yields the highest utility of all customers in its district k with c_k customers.

The probability that a district is being selected utility-dependent only depends on the customer with the highest shared utility component $u_s(j)$ in this district and the individual utility $u_i(i, k)$ of the district but is ignorant about the number of customers in this district and all other customers' shared utility component.

Bearing that in mind I define the utilization fraction f as follows. The probability of being the customer with the highest utility is $B_1(j, c_k)$ and the average number of agents driving to a district k containing customer j is $\varphi \Pi'(j)$. All agents drive to their preferred customer ($\forall i, j : d(i, j) = p(i, j)$).

$$f = \frac{1}{N} \cdot \sum_{j=1}^N \sum_{c_k=1}^j B_1(j, c_k) \cdot \left(c_k - \sum_{p_k=0}^{c_k-1} Pref(p_k, \varphi \cdot \Pi'(j)) \cdot (c_k - p_k) \right) \quad (83)$$

The utilization fraction of agents using the RD strategy is, thus, $f = f_{RD} = 30.6\%$. The expected average utility $u = 0.240 \cdot u_{max}$ is given by adapting equation 83 with the expected utility $u_e(j, k)$ for all successful agents.

$$u = \frac{1}{N} \sum_{j \in J} \sum_{c_k=1}^j B_1(j, c_k) \cdot \left(c_k - \sum_{p_k=0}^{c_k-1} Pref(p_k, \varphi \Pi'(j)) \cdot (c_k - p_k) \cdot u_e(j, k) \right) \quad (84)$$

8.5. Limited Learning

Using the strategy LL, agents first drive to the district a randomly selected customer is located in. Agents who carried a customer at time t drive to their preferred customer at time $t + 1$. The utilization fraction for the MPMC model is calculated as follows: The left summand comprises those agents who were successfully carrying a randomly chosen customer in the previous iteration (f_{t-1}) and now drive to their preferred resource. These agents are successful with probability f_{RD} . f_{RD} is the utilization fraction of the RD strategy and thus the number of customers who are preferred by any agent. The right summand comprises all other agents driving to the remaining districts. The average number of customers per district is adapted to $\lambda = \varphi(1 - f_{t-1})$ as the expected number of remaining customers is reduced.

$$f_t = f_{t-1} \cdot f_{RD} + \frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda) \cdot (c_k - o_k) \right) \quad (85)$$

For the expected utility one has to differentiate between randomly choosing agents and those who return to their preferred district, as both groups are comprised in the right summand of equation 85. Of these agents, a fraction of $\bar{r} = 0.186$ return to their preferred district, $1 - \bar{r}$ choose randomly.

$$u = f \cdot f_{RD} \cdot u_m + \frac{1}{\varphi} \cdot \sum_{c_k=1}^N C(c_k) \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda) \cdot (c_k - o_k) \right) \cdot (\bar{r} \cdot u_m + (1 - \bar{r}) \cdot u_{avg}) \quad (86)$$

From equations 85 and 86 I derive $f = 57.0\%$ and $u = 0.357 \cdot u_{max}$.

8.6. One Period Repetition

Agents adopting the OPR strategy randomly choose a resource at time t , and return there at time $t + 1$ if they were successful at time t . At time $t + 2$, agents drive to their preferred customer (after being successful at time t and $t + 1$).

The utilization fraction and utility are calculated as follows. $x = (1 - 2x) \cdot f_{NL}$ is the fraction of agents who return to the same district and who improve by driving to their preferred resource after returning to a random resource. $1 - 2x$ agents randomly select any customer, $x = (1 - 2x) \cdot f_{NL}$ of these agents are successful. f_{NL}

is the utilization fraction of the NL strategy and, therefore, randomly behaving agents. Further, all districts are utilized up to $\min(c_k, p_k)$, which comprises f_{RD} . These $f_{RD} \cdot N$ agents constantly remain with their preferred district (f_{RD} is the utilization fraction of the RD strategy).

$$f = (x + (1 - 2x) \cdot f_{NL}) \cdot (1 - f_{RD}) + f_{RD} \quad (87)$$

$$u = (x + (1 - 2x) \cdot f_{NL}) \cdot (1 - f_{RD}) \cdot u_{avg} + u_m \cdot f_{RD} \quad (88)$$

With $f_{NL} = 0.830$, and $f_{RD} = 0.308$ this results in $f = 73.9\%$ and $u = 0.457 \cdot u_{max}$.

8.7. Crowd Avoiding

The CA strategy ignores the utility or "rank" of customers; agents only drive to customers who were not carried in the previous iteration. On average, agents choose from of $\lambda = \frac{1}{1-f}$ of all customers, resulting in $\lambda \cdot c_k$ customers remaining per district.

$$f = (1 - f) \cdot \left(\sum_{c_k=1}^N C(c_k) \cdot \left(c_k - \sum_{o_k=0}^{c_k-1} O(o_k, \lambda \cdot c_k) \cdot (c_k - o_k) \right) \right) \quad (89)$$

I, therefore, conclude that the utilization fraction is $f = 49.7\%$, and that the utility is $u = 0.249 \cdot u_{max}$.

8.8. Stochastic Crowd Avoiding

Agents applying the SCA strategy either return to the same resource in the next iteration or divert to other resources. An agent i remains at district k , if k 's capacity is not fully used ($o_k \leq c_k$), or with probability $\frac{c_k}{o_k}$. If an agent i does not return to the same district, he randomly selects any resource $k \in K$.

Simulations suggest a utilization fraction of $\bar{f} = 93.8\%$ and a utility of $u = 0.469 \cdot u_{max}$.

8.9. Stochastic Rank Dependent Choice

The strategy SRD dictates that agents stochastically either drive to their preferred district or any other district, depending on the number of agents with the same preference (p_k for $p(i, k) = 1$). Diverting agents drive to (1) any underutilized district, (2) any other district, (3) the district yielding second highest utility, or (4) the underutilized district that yields the highest utility.

The overall utilization fraction f for every strategy is calculated as a generalization of equation 18.

The utilization fraction sums up the expected number of agents carrying a customer ($F(c_k, p_k)$) for the number of agents preferring district k (p_k with probability

$Pref(p_k, \varphi\Pi'(j))$, the capacity of this district (c_k with probability $C(k)$), and the customer yielding highest utility j . The utility function u analogously sums up all individual utilities $U(c_k, p_k)$ analogously.

$$f = \sum_{j \in J} \sum_{c_k=0}^N C(c_k) \cdot \sum_{p_k=0}^N Pref(p_k, \varphi\Pi'(j)) \cdot F(c_k, p_k) \quad (90)$$

$$u = \sum_{j \in J} \sum_{c_k=0}^N C(c_k) \cdot \sum_{p_k=0}^N Pref(p_k, \varphi\Pi'(j)) \cdot U(c_k, p_k) \quad (91)$$

The utilization function $F(c_k, p_k)$ is p_k , if the capacity is not exceeded by those agents preferring district k . In this case, no agent diverts and thus all agents can carry a preferred customer. Otherwise, one sums up the utilization retrieved from r_k agents redirecting for all $r_k \leq p_k$ weighted by the probability $D(p_k, c_k, r_k)$ that r_k agents divert in a district k containing c_k customers that is preferred by p_k agents and is calculated as a Poisson distribution around $p_k - c_k$ ($D(c_k, p_k, r_k) = \frac{(p_k - c_k)^{r_k}}{r_k!} \cdot e^{c_k - p_k}$). $\min(c_k, p_k - r_k)$ agents remaining at district k carry a customer in this district. If less agents divert than required, not all of them will be able to carry a customer, but all c_k customers will be carried. If more agents divert than required, all $p_k - r_k$ agents carry a customer, but not all customers are carried. Those r_k agents who redirect to another district can increase the utilization, if they are able to carry the customer they divert to. The probability of carrying a customer as a diverting agent is given by success rate s . SRD2 and SRD3 allow diverting agents to drive to fully capacitated districts. Yet, for calculating the utilization fraction I assume without loss of generality that not diverting agents favorably carry customers. Diverting agents receive a certain utilization depending on the success rate s which varies depending on the strategy and its associated behavior in case of swapping. The success rate factors in that diverting agents can only be successful if no other agent is "bullied out" his preferred district.

$$F(c_k, p_k) = \begin{cases} p_k & \text{if } p_k \leq c_k \\ \sum_{r_k=0}^{c_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k + c_k) & \\ + \sum_{r_k=c_k+1}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k + (p_k - r_k)) & \text{otherwise} \end{cases} \quad (92)$$

The utility function $U(c_k, p_k)$ is given by adapting equation 92 to cater for varying utility levels. Agents

carrying a customer from their preferred district receive on average a utility of u_m (from section 8.2.5), diverting agents receive on average u_{alt} if they are successful. u_{alt} depends on the strategy.

$$U(c_k, p_k) = \begin{cases} p_k \cdot u_m & \text{if } p_k \leq c_k \\ \sum_{r_k=0}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k \cdot u_{alt} + c_k \cdot u_m) & \\ + \sum_{r_k=0}^{p_k} D(c_k, p_k, r_k) \cdot (s \cdot r_k \cdot u_{alt} + (p_k - r_k) \cdot u_m) & \text{otherwise} \end{cases} \quad (93)$$

Table 9 compares the variables s , u_m , and u_{alt} for strategies SRD1 and SRD2. Strategies SRD3 and SRD4 perform worse than random, as first preference and alternative choice are not independent of each other (thus, diverting agents r_k are not uniformly distributed, making it impossible to analytically derive a success rate s). In simulations, the utilization fraction of SRD3 is $\bar{f} = 36.7\%$, and its utility is $u = 0.283 \cdot u_{max}$. The utilization of strategy SRD4 is $\bar{f} = 47.4\%$, and its utility is $u = 0.366 \cdot u_{max}$.

In strategy SRD1, I assume that the success rate is $s = 0.866$ as given by equation 94. On average $0.697N = (1 - 0.303)N$ agents divert to other districts. Thus, $0.697N$ customers are not being carried by an agent to whom they are first preference. I furthermore assume that these customers are Gaussian distributed across all districts, resulting in on average $\lambda = \varphi \cdot 0.697$ customers per district. The success rate s is calculated as the utilization fraction of the NL strategy with a reduced number of customers per district.

$$s = \sum_{c_k=1}^N \frac{\lambda^{c_k}}{c_k!} \cdot e^{-\lambda} \sum_{o_k=0}^{c_k-1} \frac{c_k^{o_k}}{o_k!} \cdot e^{-c_k} \quad (94)$$

In strategy SRD2, the success rate is $s = 0.850$. In this case, I calculate the expected number of previously not carried customers ($c'_k = c_k - o_k + r_k$ with probability $P(c'_k)$) and the probability that these customers are carried by r'_k agents who divert to district k .

$$s = \sum_{c_k=1}^N \sum_{c'_k}^{c_k} P(c'_k) \left(c'_k - \sum_{r'_k=0}^{c'_k-1} \frac{((1 - f_{RD}) \cdot c_k)^{r'_k}}{r'_k!} \cdot e^{-((1 - f_{RD}) \cdot c_k)} \right) \quad (95)$$

The utility u_m is the utility of strategy RD for those who are successful. I set $u_{alt} = u_{avg}$, as the alternate choice is independent from the actual utility.

Strategy	s	u_m	u_{alt}	f	u
SRD1	0.866	0.79	0.50	78.5%	0.438
SRD2	0.850	0.72	0.50	77.3%	0.432

Table 9: MPMC: SRD Strategy – Variables

8.10. Results

Table 10 shows utilization and utility for the previously examined strategies in the MPMC model.

Of the two baseline comparisons, NL outperforms RD both with respect to utilization and utility, as the number of districts containing a preferred customer is lower than a random selection of districts. None of the rank dependent strategies (LL, OPR, SRD1-SRD4) reach the utilization of the NL strategy, but OPR, SRD1 and SRD2 outperform NL with respect to utility. SCA performs best both with respect to utilization and utility.

9. Critical Discussion

In the previous sections I observe that utilization fraction and utility of some strategy vastly depend on the model variant: In general, one can state that using districts (IPMC, MPMC) improves both optimization criteria. Obviously, if there was only a single district ($D = 1, \varphi = N$) in which all customers are located, one can expect a utilization fraction of $f = 1$ regardless of the implemented strategy, as all agents can divert to other customers in the same district until every customer is carried. If there are no districts, the utilization fraction is determined by the KPRP, or the IP and MP model variant, depending on the other assumptions. I thus advise “clustering” the resources (customers) based on proximity, for example by using taxi stands. They allow agents to serve another customer in the same district if another agent already carries the selected customer. I notice that all strategies always perform at least as good in IP and IPMC as in their mixed preferences counterpart. Obviously, NL, CA and SCA are not affected, as agents never deterministically drive to their preferred resource, but utilization fraction and average agent utility for the other strategies decrease when introducing mixed preferences as the number of distinct highest utility resources decreases. The number of distinct highest utility resources depends on the probability that a resource is preferred by any given customer which is not identical for all resources in the MP and MPMC model variant but depends on the shared utility component. Due to this, exceeding f_{NL} with rank dependent strategies becomes difficult for $\alpha = 0.5$. With increasing α the number of distinct highest utility resources decreases, resulting in a decreasing utility of all rank-dependent strategies, as shown in appendix ???. Thus, I conclude that high individual utility components are preferred by

agents, as the probability of being able to carry the preferred customer increases. In mobility markets – that is vehicle for hire markets – I derive that one would prefer a high influence of the cost of driving to the pickup location which can either be achieved by revenue in a small range or by high distances to the pickup location. Alternatively, a coordination instance could impose personalized incentives, causing agents to distribute themselves in balance with customers.

I also observe that stochastic rank dependent strategies (SRD) outperform their strict counterpart (RD). This is because a fraction of agents chooses its top preference, whilst the other agents can receive utility from another resource. I observe that SRD1 (and SRD4 in IP and IPMC) perform best with respect to utilization fraction f (most customers are carried). SRD1 and SRD4 outperform SRD2 and SRD3 in the IP and IPMC model variants of the VFHP, and SRD1 outperforms SRD2 in the MP and MPMC model variants, as the success rate of redirecting agents is higher. In IP and IPMC, SRD4 outperforms SRD1 with respect to utility, as agents always choose a district yielding high utility. SRD4 performs poorly for mixed utility models (MP, MPMC), as most agents share the same highest utility district with remaining utility. Yet, SRD2 and SRD3 require less information about the preferences of other agents and are therefore preferred in environments without full information.

The CA strategy outperforms the NL strategy in none of the models and is more complex as it requires information about the occupancy rate of all resources, making it unsuitable for implementation. The LL strategy is outperformed by the OPR strategy in all models, making it less attractive for implementation. Yet, the two-step approach is easier to establish in a larger group of agents. From comparing the strategies LL and OPR I conclude that waiting for m periods before choosing the highest utility resource further improves both optimization criteria (strategy m-Period Repetition, mPR). I observe that OPR and SCA perform best regarding the utilization fraction and utility. Yet, agents will not be able to carry their top priority customer with SCA in most cases (probability $\frac{1}{N}$). My findings recommend that taxi drivers consider both history and associated utility when choosing a customer or resource.

Yet, my model draws a rather theoretical picture of the reality: I assume that utilities $u_s(j)$ and $u_i(i,j)$ are uniformly distributed and random ($\frac{1}{N} \dots 1$ with step size

Strategy	utilization f	utility u
NL	83.0%	0.415
RD	30.6%	0.240
LL	57.0%	0.357
OPR	73.9%	0.457
CA	49.7%	0.249
SCA	93.8%	0.469
SRD1	78.5%	0.438
SRD2	77.3%	0.432
SRD3	36.7%	0.283
SRD4	47.4%	0.366

Table 10: MPMC: Comparing Strategies

$\frac{1}{N}$), allowing for an analytical approach. In most cities, one would rather assume a majority of customers returning a low or medium utility and only very few trips with very high utility. Also, assuming Gaussian-distributed numbers of customer per district is a major abstraction, in reality, a small number of hot spots such as airports or railway stations draw more attention than a large number of residential neighborhoods. Yet, the VFHP game model I discussed in chapters 5-8 can easily be adapted by exchanging $C(c_k)$ by more suitable functions for the given distribution. In the MP and MPMC model variants, I model the distance between agent i and customer j as $u_i(i, j)$. In real world examples, $u_i(i, j)$ depends on the history, as agents move through the city. Also, two adjacent resources will result in similar utilities for all agents which is not reflected in the presented model. Though, my model allows for extensions addressing these limitations.

In reality, the individual utility of agents – that is distance between agent and resource – changes in every iteration, as agents drive to customers. Thus, the utility agents can derive from customers has to be recalculated in every iteration. Yet, varying utilities do not influence the general idea VFHP game model; one only had to retrieve information about the preferences of all other agents in every iteration. Another abstraction concerns the timing between agents: One cannot assume that all agents select a resource at the same time. One could impose a discrete time model assuming that every agent drives to one customer per discrete time step, but as driving to a customer takes differently long depending on the distance. In the VFHP game model, it is sufficient to assume that the number of customers and agents is identical in all iterations, but several of the history-dependent strategies (LL, OPR, SCA) will perform differently for agents who did not participate in the previous iteration, as these agents will have to select a random resource rather than using a more promising selection. For example, agents implementing the OPR strategy receive a certain utilization of $f(i) = 1$ from customer j in the “return” phase, as no other agent

drives to this customer j if this resource was occupied in the previous iteration. Yet, if agent i returns to a customer after pausing for several iterations, it is possible that another agent chose this resource as well, reducing utilization fraction and utility. Also, drivers who did not carry a customer will be able to drive to another customer directly after, whilst agents carrying a customer first have to finish this trip and are thus not available during the next iteration. One can extend the VFHP game model with a “continue carrying” phase for agents, in which they are utilized ($f(i) = 1$) and the utility the carried customer yields is divided up over the all iterations this trip takes. Customers disappear after being carried, and new customers appear frequently. As the shared utility of customers is the expected revenue, the VFHP game model can easily incorporate appearing and disappearing customers. Also, the expected utility yielded by customers can be difficult to determine, as individual behavior cannot be predicted precisely. It is possible to predict general tendencies (e.g., customers at airports often travel downtown and thus quite far), but for other locations, one cannot predict precise travel distances or patterns of customers (e.g. in city centers, most customers travel short distance, but few customers need longer transport, yet, it is difficult to predict when exactly customers require these longer trips). The IP and IPMC model variant do not use shared utilities in terms of customer revenue and are therefore more suited if the utility is unknown. In more rural areas, the expected number of customers in a district can be below 1, but the VFHP assumes discrete numbers of customers per district. Whilst rounding is reasonable for larger numbers of customers per district, rounding will frequently result in no expected customers in rural areas. There, vehicles for hire are usually called by phone. Thus, a dispatcher sends a driver to pick up this customer. The VFHP on the opposite mimics taxi hailing or calling a nearby taxi via app, if no dispatcher is available.

Despite the above limitations, the VFHP presents a

suitable game model for agent behavior in vehicle for hire markets and lays ground work for improving utilization and utility in mobility markets.

10. Conclusion and Future Work

In this thesis I analysed two different models for mobility markets, the Kolkata Paise Restaurant Problem (KPRP) and four model variants of the Vehicle for Hire Problem (VFHP). To adapt the KPRP for mobility markets, I gradually drop or alter the assumptions of the KPRP: Agents no longer agree upon the resources' utilities (IP and MP model variants), and resources are "clustered" in districts, allowing agents to deviate from their first choice (IPMC and MPMC model variants). Further, I compared those five models by testing for utilization fraction and utility for agents using one of seven different strategies. Three of these strategies stem from Chakrabarti et al. (2009), two further strategies were introduced by Ghosh et al. (2013). I developed the strategies RD and SRD to specifically address the requirements of dynamic mobility markets. In dynamic matching markets, the behavior of other agents in previous iterations cannot determine the utility agents associate with resources in the future with absolute certainty as agents and customers enter and leave the market at will, calling for history-independent rank-dependent strategies.

Future research will be conducted on (1) behavior of agents, if two or more strategies are implemented in one market and the influence on utilization fraction and utility, (2) performance of the discussed strategies in practice, (3) incentive mechanisms and their effect in practice, and (4) the influence of the rise of autonomous cars and successive merge of the vehicle for hire and the car-sharing market.

If agents apply different strategies, the overall utilization fraction and utility might increase or decrease. Also, the utility could be unevenly distributed. For example, if $N - 1$ agents play NL in the KPRP and one agent plays RD, this agent can expect a higher utility than the other agents ($0.632 \cdot u_{max}$ vs. $0.316 \cdot u_{max}$). Unilateral deviation can therefore be beneficial for agents. In the CA strategy, unilaterally deviating agents can implement a strategy in which they only choose from previously occupied resources, if only one agent deviates, he is guaranteed a utilization of $f(i) = 1$. The OPR strategy retrieves its high utility from agents not randomly choosing resources which were served by other agents the previous iteration, including those agents who constantly carry their preferred customer. Single agents implementing a NL strategy reduce the number of agents returning to their preferred resource, decreasing the performance of the OPR strategy.

This thesis focuses on the performance of several strategies in theoretical settings. As discussed in chapter 9, utilization and utility can vary as the assumptions of the VFHP deviate from reality. With real world data on the location of customers during a given time frame and the routes of drivers, one can evaluate whether the strategies improve current driver behavior. With insight from this data analysis, one can improve the strategies presented in this thesis and continue with incentive mechanisms to enforce beneficial behavior.

One can use the knowledge about the theoretic (and real world) performance of different strategies to incentivize behavior that is beneficial for the entire group. As discussed in chapter 9, agents incorporating the strategy OPR achieve a high utilization fraction and high utility in the IP and IPMC model. The strategy dictates a three-step approach: A random choice of a resource, returning to this resource once, and driving to the preferred resource. Yet, agents might be reluctant to wait for one iteration prior to driving to the preferred resource (e.g. due to missing trust in other agents, bounded rationality). For these agents, a coordination instance can offer incentives to return to the same resource.

Developments in the field of autonomous cars will most likely result in the end of vehicle for hire markets in its current setup, as drivers are no longer required, but cars independently carry passengers. Another industry that develops towards autonomous vehicles for passenger transportation is the car-sharing market in which passengers can rent cars for a short period (i.e. for one-way trips in major cities). The vehicle for hire market and car-sharing market steer towards offering the same service, if drivers become obsolete. Obviously, strategies and algorithms to redirect agents will become increasingly important; future research should therefore focus on improving the basic strategies presented in this thesis.

References

- Abraham, D. J., Irving, R. W., Kavitha, T., and Mehlhorn, K. Popular matchings. *SIAM Journal on Computing*, 37(4):1030–1045, 2007.
- Agarwal, S., Ghosh, D., and Chakrabarti, A. S. Self-organization in a distributed coordination game through heuristic rules. *arXiv preprint arXiv:1608.00213*, 2016.
- Akbarpour, M., Li, S., and Oveis Gharan, S. Thickness and information in dynamic matching markets. Available at SSRN 2394319, 2016.
- Alonso-Mora, J., Samaranyake, S., Wallar, A., Frazzoli, E., and Rus, D. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, page 201611675, 2017.
- Arnott, R. Taxi travel should be subsidized. *Journal of Urban Economics*, 40(3):316–333, 1996.
- Arthur, W. B. Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406–411, 1994.
- Bloch, F. and Houy, N. Optimal assignment of durable objects to successive agents. *Economic Theory*, 51(1):13–33, 2012.
- Cairns, R. D. and Liston-Heyes, C. Competition and regulation in the taxi industry. *Journal of Public Economics*, 59(1):1–15, 1996.
- Chakrabarti, A. S., Chakrabarti, B. K., Chatterjee, A., and Mitra, M. The kolkata paise restaurant problem and resource utilization. *Physica A: Statistical Mechanics and its Applications*, 388(12):2420–2426, 2009.
- Challet, D. and Zhang, Y.-C. On the minority game: Analytical and numerical studies. *Physica A: Statistical Mechanics and its Applications*, 256(3):514–532, 1998.
- Chen, M. K. and Sheldon, M. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. Technical report, Mimeo, UCLA, 2015.
- Chen, Y. and Hu, M. Pricing and matching with forward-looking buyers and sellers. 2016.
- Chmura, T. and Pitz, T. Successful strategies in repeated minority games. *Physica A: Statistical Mechanics and its Applications*, 363(2):477–480, 2006.
- Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. On the inefficiency of equilibria in congestion games. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 167–181. Springer, 2005.
- Cramer, J. and Krueger, A. Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–182, 2016.
- Gale, D. and Shapley, L. S. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- Ge, Y., Xiong, H., Tuzhilin, A., Xiao, K., Gruteser, M., and Pazzani, M. An energy-efficient mobile recommender system. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 899–908. ACM, 2010.
- Ghosh, A., Biswas, S., Chatterjee, A., Chakrabarti, A. S., Naskar, T., Mitra, M., and Chakrabarti, B. K. Kolkata paise restaurant problem: An introduction. In *Econophysics of Systemic Risk and Network Dynamics*, pages 173–200. Springer, 2013.
- Hall, J., Kendrick, C., and Nosko, C. The effects of uber's surge pricing: A case study. *The University of Chicago Booth School of Business*, 2015.
- Kim, H., Yang, I., and Choi, K. An agent-based simulation model for analyzing the impact of asymmetric passenger demand on taxi service. *KSCE Journal of Civil Engineering*, 15(1):187–195, 2011.
- Kurino, M. House allocation with overlapping agents: A dynamic mechanism design approach. Technical report, Jena economic research papers, 2009.
- Lasry, J.-M. and Lions, P.-L. Mean field games. *Japanese journal of mathematics*, 2(1):229–260, 2007.
- Lee, D.-H., Wang, H., Cheu, R., and Teo, S. Taxi dispatch system based on current demands and real-time traffic conditions. *Transportation Research Record: Journal of the Transportation Research Board*, (1882):193–200, 2004.
- Leshno, J. D. Dynamic matching in overloaded systems. URL <http://www.people.fas.harvard.edu/jleshto/papers.html>, 2:7, 2012.
- Li, B., Zhang, D., Sun, L., Chen, C., Li, S., Qi, G., and Yang, Q. Hunting or waiting? discovering passenger-finding strategies from a large-scale real-world taxi dataset. In *Pervasive Computing and Communications Workshops (PERCOM Workshops)*, 2011 IEEE International Conference on, pages 63–68. IEEE, 2011.
- Li, S. S. H. *Multi-attribute taxi logistics optimization*. PhD thesis, Massachusetts Institute of Technology, 2006.
- Linne+Krause Marketing-Forschung. Gutachten über die wirtschaftliche lage des hamburger taxi-gewerbes. http://hh-taxi.de/wp-content/Taxi_Gutachten_L&K_2011.pdf, 2011. retrieved at 12.02.2017.
- Linne+Krause Marketing-Forschung. Untersuchung zur wirtschaftlichkeit des berliner taxigewerbes. http://www.stadtentwicklung.berlin.de/verkehr/politik/taxi/download/untersuchung_wirtschaftlichkeit_taxi_berlin.pdf, 2016. retrieved at 12.02.2017.
- Manlove, D. F. and Sng, C. T. Popular matchings in the capacitated house allocation problem. In *European Symposium on Algorithms*, pages 492–503. Springer, 2006.
- Milchtaich, I. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.
- Monderer, D. and Shapley, L. S. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- Nash, J. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- Rayle, L., Shaheen, S., Chan, N., Dai, D., and Cervero, R. App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco. *University of California, Berkeley, United States Rogers, B.(2015) The social costs of Uber. James E. Beasley School of Law, Temple University, Philadelphia, United States*, 2014.
- Rogers, B. The social costs of uber. *U. Chi. L. Rev. Dialogue*, 82:85, 2015.
- Rosenthal, R. W. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- Seow, K. T., Dang, N. H., and Lee, D.-H. A collaborative multiagent taxi-dispatch system. *IEEE Transactions on Automation Science and Engineering*, 7(3):607–616, 2010.
- Shi, Y. and Lian, Z. Optimization and strategic behavior in a passenger-taxi service system. *European Journal of Operational Research*, 249(3):1024–1032, 2016.
- Trigo, P., Jonsson, A., and Coelho, H. Coordination with collective and individual decisions. In *Advances in Artificial Intelligence-IBERAMIA-SBIA 2006*, pages 37–47. Springer, 2006.
- Wong, R., Szeto, W., and Wong, S. A two-stage approach to modeling vacant taxi movements. *Transportation Research Procedia*, 7:254–275, 2015.
- Yang, H., Ye, M., Tang, W. H.-C., and Wong, S. C. A multiperiod dynamic model of taxi services with endogenous service intensity. *Operations research*, 53(3):501–515, 2005.
- Yang, P., Frazier, P. I., et al. Mean field equilibria for competitive exploration in resource sharing settings. In *Proceedings of the 25th International Conference on World Wide Web*, pages 177–187. International World Wide Web Conferences Steering Committee, 2016.
- Ye, Q. C. and Zhang, Y. Participation behavior and social welfare in repeated task allocations. In *Agents (ICA)*, IEEE International Conference on, pages 94–97. IEEE, 2016.
- Ye, Q. C., Zhang, Y., and Dekker, R. Fair task allocation in transportation. *Omega*, 68:1–16, 2017.
- Yuan, J., Zheng, Y., Zhang, L., Xie, X., and Sun, G. Where to find my next passenger. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 109–118. ACM, 2011.



Word embedding, neural networks and text classification: what is the state-of-the-art?

Estevan Vilar

ESCP Europe

Abstract

In this bachelor thesis, I first introduce the machine learning methodology of text classification with the goal to describe the functioning of neural networks. Then, I identify and discuss the current development of Convolutional Neural Networks and Recurrent Neural Networks from a text classification perspective and compare both models. Furthermore, I introduce different techniques used to translate textual information in a language comprehensible by the computer, which ultimately serve as inputs for the models previously discussed. From there, I propose a method for the models to cope with words absent from a training corpus. This first part has also the goal to facilitate the access to the machine learning world to a broader audience than computer science students and experts.

To test the proposal, I implement and compare two state-of-the-art models and eight different word representations using pre-trained vectors on a dataset given by LogMeIn and on a common benchmark. I find that, with my configuration, Convolutional Neural Networks are easier to train and are also yielding better results. Nevertheless, I highlight that models that combine both architectures can potentially have a better performance, but need more work on identifying appropriate hyperparameters for training. Finally, I find that the efficacy of word embedding methods depends not only on the dataset but also on the model used to tackle the subsequent task. In my context, they can boost performance by up to 10.2% compared to a random initialization. However, further investigations are necessary to evaluate the value of my proposal with a corpus that contains a greater ratio of unknown relevant words.

Keywords: neural networks; machine learning; word embedding; text classification; business analytics

1. Introduction

“Innovation is hard. It really is. Because most people don’t get it. Remember, the automobile, the airplane, the telephone, these were all considered toys at their introduction because they had no constituency. They were too new.” Nolan Kay Bushnell

1.1. Data availability

Data has been called by The Economist the “new oil” (Economist, 2017) as they are now “abundant, ubiquitous and far more valuable [than before]”. Internet, social media, sensors, and smartphones have all contributed to the production of electronic information whether structured or not. Daily, 2.5 quintillion bytes of data are created (IBM, 2018). With this increasing amount of data, a need to accurately extract, integrate and classify these resources has appeared in the last two decades.

Among this electronic information, a plethora of textual resources such as tweets, reviews, comments, emails or news but also scanned documents or handwritten notes are produced, and therefore techniques in the field of Natural Language Processing (NLP) and machine learning have been developed to get meaningful knowledge from this information.

The first goal of this bachelor thesis in collaboration with the company LogMeIn Inc.¹ is to evaluate the current state-of-the-art of classification techniques with neural networks, select the appropriate algorithms and subsequently tackle the automated classification and performance analysis. These tasks will be performed to pinpoint the most effective method to sort textual reviews of customers about the use of GoToMeeting² - an online meeting, desktop sharing and video conferencing software - by subject (audio, non-audio).

Second, this work is also exploratory as a new method to deal with out-of-vocabulary words is tested and compared

¹<https://www.logmein.com/>

²<https://www.gotomeeting.com/>

with the state-of-the-art. The goal is to improve the generalization power of classification methods, without deep and heavy implementations.

1.2. Feedback loops

Part of the agile methodology, experimentation is favoured over elaborate planning and so is customer feedback over intuition (Rahimian and Ramsin, 2008). As a consequence, one component of the methodology is to enter quickly what is commonly called feedback loops. It consists of building a minimum viable product, getting customers' feedback and used it to improve the product. In that context, many online tools have been developed to conduct surveys, but also many applications such as AirBnB³ or Uber⁴ include reviews as part of their product to gain the trust of their users. If these tools allow developers and managers to collect a significant amount of data, there is, however, a need to efficiently analyse these data to perform qualitative analysis and infer where resources should be allocated. The third goal of this thesis is, therefore, to present tools that managers or entrepreneurs can leverage to build better products faster. As a consequence, this thesis has been written with a goal in mind to facilitate the access to modern tools for analysis, more specifically neural networks, to add a new card in the hands of managers to understand their customer concerns better.

In Section 2, I introduce the theoretical background and different concepts necessary to understand the functioning of neural networks. I describe two commonly used architectures namely Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) and compare their performance when it comes to classification tasks. In Section 3, I discuss the conversion of textual information in a format recognisable by computers. I introduce three techniques to extract information from texts: GloVe, Word2Vec, and Fast-Text. I also propose a method to deal with words that are not present in the training data. In Section 4, I describe the benchmark to compare the models introduced in Section 2 and techniques mentioned in Section 3. Section 5 includes the results and discussions following the experiment and Section 6 is the concluding part of this thesis.

2. Text Classification and Machine Learning

“Science is the systematic classification of experience” George Henry Lewes

Text Classification (TC) (also called text categorisation or topic spotting) refers to the identification and labelling of themes or topics of a sentence or document (Sebastiani, 2002). An example would be to label a comment based on the topic it covers like “audio”, “screen” and “video”. In the early 90's, the emergence of digital data, and the growing

computational power of machines contributed to the development of the field. Also, the broad applicability of the task in activities such as spam detection, metadata generation or organisation of documents attracted the interest of technological companies. Before that time, techniques involved knowledge engineering (KE) which consists of classifying a textual document based on knowledge encoded in a set of rules manually defined (Faraz, 2015). However, in the 90's the machine learning (ML) paradigm shifted the attention of researchers away from KE. Rather than imposing classification rules to machines, researchers started to build solutions that let the computer deduce the attributes that will lead to efficient classification. From a pre-classified set of documents, the machine would thus learn the characteristics of interests to build an automated classifier.

Formally, TC tasks assign a Boolean value to each pair $\langle d_j, c_j \rangle \in D \times C$, with D being a training set of documents and $C = \{c_1, \dots, c_n\}$ a set of predefined categories. The goal is to approximate the target function $f: D \times C \rightarrow [T, F]$ where T indicates that d_j must be classified under c_i whereas F indicates that d_j must not be classified under c_i . As f is unknown, the function $g: D \times C \rightarrow [T, F]$ that approximate f - also called classifier (or model, or rule) - is used. Then, the effectiveness of the classifier - or accuracy - refers to the degree to which f and g coincide. Ultimately, classifying a document D under $C = [c_1, \dots, c_i, \dots, c_n]$ with $i=1 \dots n$ can be seen as n independent problems with $f_i: D \rightarrow [T, F]$ as an unknown target function for c_i and $g_i: D \rightarrow [T, F]$ a classifier for c_i (Sebastiani, 2002).

The first challenge lies in the so-called inter-indexed inconsistency based on the first law of Jesse H. Shera (Cleverdon, 1984). It states that “No cataloguer will accept the work of any other cataloguer”. This law highlights the subjectivism of classification tasks and therefore points to the non-existence of a deterministic solution - a function f - for the classification problem. Nevertheless, in the last decades, researchers have been looking for an optimal function g to solve specific classification problems.

In ML, building a classifier relies on the availability of a pre-classified corpus from which to deduce the relevant characteristics i.e., a corpus on which the values of every pair $\langle d_j, c_j \rangle \in D \times C$ are known. Besides, to evaluate the effectiveness of the classifier, it is common practice to split this pre-classified corpus between a training- set - used to build the classifier - and a test set - to assess the effectiveness of the classifier. Once the classifier is built, each d_j from the test set are used as input which produces a corresponding c_i . The effectiveness is measured by how often the pairs $\langle d_j, c_i \rangle$ matches the values of the pre-classified corpus while testing.

Since the beginning of machine learning techniques for TC, a broad range of model including rule induction, naïve-bays, decisions trees, K-nearest neighbours (KNN), support vector machines (SVM) and neural networks have been used to build classifiers. A comparative study of the techniques is available in (Kaur and Kaur, 2017; Khan et al., 2010; Nikam, 2015). As pointed out in (Young et al., 2018), deep learning architecture such as deep neural networks have increasingly

³<https://www.airbnb.com/>

⁴<https://www.uber.com/en-MX/>




Figure 1: Percentage of deep learning papers in ACL, EMNLP, EACL, NAACL over the last six years; Source: (Young et al., 2018)

attracted the attention of researchers as shown in Figure 1. For that reason; this work is focusing on neural networks for TC tasks.

2.1. Neural Networks for text classification

Artificial neural networks as defined by Dr. Robert Hecht-Nielsen quoted in Neural Network Primer: Part 1 is:

“a computing system made up of a number of simple, highly interconnected processing elements, which process information by their dynamic state response to external inputs.” (Caudill, 1986).

If the essential components of neural networks remain the same, their architecture can change a lot. For this section, I aim to identify the state-of-the-art model for TC among Convolutional Neural Networks (CNN) and Recurrent Neural Network (RNN). First, I describe the functioning of a multilayer perceptron (MLP), a feed-forward neural network, which represents one of the most straightforward architectures of neural networks⁵. It is done to introduce the fundamental concepts necessary to understand the CNN and RNN. I describe the models and their most up-to-date applications for TC tasks.

2.2. Feed-forward neural networks

The definition mentioned previously encompasses the essential components of modern neural networks. Hecht-Nielsen refers to what are today called neurons with the word “processing elements”. This computational unit receives a set of scalar x_i or vector x as input, (1) multiplies

them by their importance - their weights w_i , (2) and apply a function f such as summation or max operation. Finally, (3) it applies a non-linear function g - also called activation function - on the result, which represents the output - a single scalar y or vector y as shown in Figure 2.

Artificial neural networks are made out of a multitude of neurons that are interconnected in different layers as illustrated in Figure 3.

They have the power to approximate any Borel functions from a finite dimensional space to another as shown in (Hornik et al., 1989), a category under which classifiers defined in the previous paragraph fall.

In mathematical notations, the feed-forward neural network represented in Figure 3 with two hidden layers would be expressed as follow⁶:

$$NN_2(x) = g^2(g^1(xW^1)W^2)W^3$$

with $x \in R^{input}$ an input vector (dimension of the Figure is 3), $W^1 \in R^{input} \times R^{output}$ is the weight matrix from the input to the first hidden layer, $W^2 \in R^{input} \times R^{output}$ is the weight matrix from the first hidden layer to the second hidden layer $W^3 \in R^{input} \times R^{output}$ is the weight matrix from the second layer to the output layer, $g^1()$ is the activation function in the first layer and $g^2()$ is the activation function of the second layer. In Figure 3, W^1 , W^2 , and W^3 are of dimension 3×3 , 3×2 and 2×4 respectively.

⁶Some feed-forward neural networks include a bias term in some layer, which is a neuron that is not connected with the previous layer. Figure 3 does not have any bias and therefore the bias is not included in the mathematical notation.

⁵The simplest one is a single layer perceptron




Figure 2: Illustration of the tasks performed in a neuron; Source: Author’s own representation




Figure 3: Feed-forward neural network with two hidden layers; Source: Author’s own representation

Alternatively, the hidden layers could be expressed as:

$$h^1 = g^1(xW^1) \text{ for the first layer}$$

$$h^2 = g^2(h^1W^2) \text{ for the second layer}$$

This gives us:

$$NN_2(x) = h^2W^3$$

The collection of matrices W^1, W^2, W^3 is referred in the

literature as the parameters θ of the neural network. In classification problems, feed-forward neural networks are often designed such as each element in the output layer is positive and that they sum to 1. The output vector can, therefore, be interpreted as a probability distribution over the different classes $[c_1, \dots, c_n]$. This final transformation is often performed with a softmax function⁷.

⁷softmax(x_i) = $\frac{e^{x_i}}{\sum_{j=1}^k e^{x_j}}$ for $x = x_1 \dots x_k$

Table 1: Summary of the most commonly used activation functions and their first derivative; Source: Compiled by author

Name	$f(x)$	$f'(x)$
Sigmoid	$\frac{1}{1+\exp(-x)}$	$\text{sigmoid}(x)(1 - \text{sigmoid}(x))$
Tanh	$\frac{e^x - e^{-x}}{e^x + e^{-x}}$	$\frac{d}{dx} \tanh(x) = 1 - \tanh(x)^2$
ReLU	$\max(0, x)$	$\begin{cases} 0, & x < 0 \\ 1, & x \geq 0 \end{cases}$
ELU	$\begin{cases} \alpha(\exp(x) - 1), & x < 0 \\ x, & x \geq 0 \end{cases}$	$\begin{cases} \alpha(\exp(x) - 1), & x < 0 \\ 1, & x \geq 0 \end{cases}$
Swish	$x * \text{sigmoid}(\beta x)$	$\beta \text{swich}(x) + \text{sigmoid}(\beta x)(1 - \beta(\text{swich}(x)))$

2.2.1. Input layer

The input of the neural network is usually a vector $x = (x_1, \dots, x_k)$. For TC problems, this vector is the result of a transformation of textual data to a vector representation. It is often referred as an embedding layer. I discuss the vector representation of text in paragraph 3.1.

2.2.2. Activation functions

In the machine learning literature, many activation functions including sigmoid, Rectified Linear Units (ReLU) (Hahnloser et al., 2000), Exponential Linear Unit (ELU) (Clevert et al., 2015) and tanh have been considered yielding different results. In its paper, Alcantara (2017) provides a comparison of different activation functions and concludes that ELU performs the best with ReLU nevertheless yielding great results. However, in the recent work of Ramachandran et al., 2017, the authors claimed that no other function had been more adopted than ReLU thanks to its simplicity and effectiveness. It is also concluded that Swish, a function similar to the Sigmoid-weighted Linear Unit (SiL) (Elfwing et al., 2018) performed better than ReLU. As far as I know, no comparison between Swish and ELU has been made and it is still an open-question to determine which one performs best. A summary of common activation functions is available in Table 1. Also, plots of these functions are available in Figure 4 and Figure 5.

2.2.3. Training a neural network

Neural networks must be trained to be efficient. Training a neural network involves setting the right weights in the various matrices W : in another word, tuning the parameters θ the best possible so the neural network approximates the desired function. To do so, a loss function is optimised and various techniques to perform the task exist.

Loss functions

As mentioned in paragraph 2.1, it is common practice in machine learning to split the data into a training set and a validation one. Let us define the output of the neural network as \hat{c} and the actual output as c . In the training set, all the pairs $\langle d_j, c_j \rangle \in D \times C$ - each document and their corresponding class c_i - are known. The objective of the training is to

minimise the function $L(\hat{c}, c)$ - a loss function - that gives a score to \hat{c} based on c . The score is therefore null if $\hat{c}_i = c_i$ and positive otherwise.

Recently, a comparison between several loss functions has been performed for TC purpose (Janocha and Czarnecki, 2017). Out of 12 loss functions - showed in Table 2 - the authors conclude that non-log losses are preferable for classification purpose. In particular, they identify the squared hinge loss (formula present in Table 1) to be the best performing function. They note however that if much noise⁸ is present in the data set, the expectation loss is the preferable choice.

Training techniques

The loss function is what needs to be minimised, and the computer must be told how to do it i.e., defining a training algorithm for the neural network.

As pointed out in (Goodfellow et al., 2016, Chapter 8), the training of the parameters θ is indirect as we hope by minimising $L(\hat{c}, c)$ we will obtain the best parameters. Therefore the techniques differ from classic optimisation problems. This include for example not evaluating the loss on the whole data set but rather on small batches and then average the results for computation power purpose. Indeed, as the standard error of the mean from a sample n is $\frac{\sigma}{\sqrt{n}}$ where σ is the true standard error, training a set of 10'000 examples takes 100 times more computational power than training a set of 100 examples, but reduces the error only by a factor 10 (Goodfellow et al., 2016, Chapter 8). Using less than all the training examples available is referred as mini-batch methods.

The most used category of optimisation algorithm are named back propagation or backward propagation of errors (Rumelhart et al., 1988) and its best representative is currently the stochastic gradient descent (SGD) (Goodfellow et al., 2016; Ruder, 2016). It consists of an iterative approach that reduces $L(\hat{c}, c)$ by moving the parameters θ in the direction opposite to sign of $L'(\hat{c}, c)$ - the derivative of the loss function. The algorithm is shown in Figure 6.

The learning rate ϵ_k present in Figure ?? as a required output is a parameter that defines how quickly the old pa-

⁸i.e. a big variability




Figure 4: Plots of activation functions including Sigmoid, ELU, ReLU; Source: (Alcantara, 2017)




Figure 5: Plot of the Swish function with different betas; Source: (Ramachandran et al., 2017)

Algorithm Stochastic gradient descent (SGD) update at training iteration k

Require: Learning rate ϵ_k .

Require: Initial parameter θ

while stopping criterion not met **do**

Sample a minibatch of m examples from the training set $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ with corresponding targets c_i .

Compute gradient estimate: $\hat{g} \leftarrow +\frac{1}{m} \nabla_{\theta} \sum_i L(g(\mathbf{x}_i; \theta), c_i)$

Apply update: $\theta \leftarrow \theta - \epsilon \hat{g}$

end while

Figure 6: Algorithm of SGD. Note that the function $g(x_i; \theta) = \hat{c}_i$ is the one referred to in 2; Source: (Goodfellow et al., 2016, Chapter 8)

parameters are forgotten compared to the new one. It has been demonstrated that, if the learning rate is appropriately set, using SGD, the function will surely converge to a global minimum or local minimum, if the function is convex (such as the

one presented in 2.2.3.1) (Bottou, 1998; Kiwiel, 2001). Furthermore, Bottou, 2012 suggest to update the learning rate in function of the iteration - also called epoch - as follow:

Table 2: List of loss functions tested in (Janocha and Czarnecki, 2017). The authors name “y” the true value. I use the notation c. Similarly, the output of the neural network is named “o” whereas I name it \hat{c} ; Source: (Janocha and Czarnecki, 2017)

symbol	name	equation
Λ_1	L_1 loss	$\ y - o\ _1$
Λ_2	L_2 loss	$\ y - o\ _2^2$
$\Lambda_1 \circ \sigma$	expectation loss	$\ y - \sigma(o)\ _1$
$\Lambda_2 \circ \sigma$	regularised expectation loss	$\ y - \sigma(o)\ _2^2$
$\Lambda_\infty \circ \sigma$	Chebyshev loss	$\max_j \sigma(o)^{(j)} - y^{(j)} $
<i>hinge</i>	hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{y}^{(j)} o^{(j)})$
<i>hinge</i> ²	squared hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{y}^{(j)} o^{(j)})^2$
<i>hinge</i> ³	cubed hinge (margin) loss	$\sum_j \max(0, \frac{1}{2} - \hat{y}^{(j)} o^{(j)})^3$
<i>log</i>	log (cross entropy) loss	$-\sum_j y^{(j)} \log \sigma(o)^{(j)}$
<i>log</i> ²	squared log loss	$-\sum_j [y^{(j)} \log \sigma(o)^{(j)}]^2$
<i>tan</i>	Tanimoto loss	$\frac{-\sum_j \sigma(o)^{(j)} y^{(j)}}{\ \sigma(o)\ _2^2 + \ y\ _2^2 - \sum_j \sigma(o)^{(j)} y^{(j)}}$
D_{CS}	Cauchy-Schwarz Divergence	$-\log \frac{\sum_j \sigma(o)^{(j)} y^{(j)}}{\ \sigma(o)\ _2 \ y\ _2}$

$$\epsilon_k = \epsilon_0 \frac{1}{1 + \epsilon_0 \delta k}$$

With ϵ_0 the initial learning rate and δ a hyperparameter⁹ to be set. However, as pointed out in (Zeiler, 2012), setting the hyperparameters alter the results of the neural networks, and the tuning can be tricky. He, therefore, presents an improvement of the standard SGD, ADADELTA, that, when used, the performance of the neural networks is not sensitive on the hyperparameter of the learning rate. The algorithm is shown in Figure 7.

Moreover, similar algorithms to ADADELTA exist such as ADAM (Kingma and Ba, 2015) or Nadam (Dozat, 2016). Finally, it is worth to point out that Ranganathan and Nataraian (2018) recently developed a new method of backpropagation without using SGD but rather Moore-Penrose Pseudo Inverse¹⁰ with promising results.

Initialisation of the network

At the beginning of the training, the weights in the different matrices W must be set. This point can determine whether the loss function - regardless of its form - will converge or diverge. Therefore, two underlying questions emerge from this issue: what is the ideal magnitude of the initial weights and what is the range in which they must be included? Before

2006, deep neural networks tended to produce inaccurate results and one reason for that is that initialisation of the network was usually totally random (Erhan et al., 2009; Glorot and Bengio, 2010; Sutskever et al., 2013). This resulted in errors such as vanishing (converging close to 0) or exploding (becoming high) gradients which does not allow the neural network to approximate the required function. In addition, neurons tended to become saturated - setting output value to 0 due to very small gradients. Likewise, output values could become too high - or die - resulting in a gradient of 0 due to inputs being negative caused by a big negative change in the gradient during the previous iteration. These issues can be solved with a wise choice of the loss function, learning algorithm, and effective initialisation of the network. An initialisation method - the xavier initialisation - introduced in (Glorot and Bengio, 2010), has become a popular technique among researchers (Goldberg, 2015). It consists of initialising the matrix as follow:

$$W \sim U\left[-\frac{\sqrt{6}}{\sqrt{d_{in} + d_{out}}}; \frac{\sqrt{6}}{\sqrt{d_{in} + d_{out}}}\right]$$

With $U[a, b]$ being a uniformly sampled value between a and b , d_{in} is the dimension of the input vector, and d_{out} is the dimension of the output vector. Using this initialisation makes sure that the distribution of the input is centred around 0 and of variance 1. However, this method assumes that the activation function is linear which is not the case for ReLu for instance. Also, this method seems not to work for very deep models (Glorot and Bengio, 2010). He et al.

⁹In machine learning, the word “hyperparameter” is used to distinguish from the parameters θ . Hyperparameters are higher level parameters set to configure properties of the neural network.

¹⁰A generalisation of the notion of inverse matrix that satisfies the four Moore-Penrose conditions (Penrose, 1955)

Algorithm Computing ADADELTA update at time t

Require: Decay rate ρ , Constant ϵ
Require: Initial parameter θ
Initialize accumulation variables $E[g^2]_0 = 0, E[\Delta\theta^2]_0 = 0$
for $t = 1 : T$ **do** %% Loop over # of updates
 Compute Gradient: g_t
 Accumulate Gradient: $E[g^2]_t = \rho E[g^2]_{t-1} + (1 - \rho)g_t^2$
 Compute Update: $\Delta x_t = -\frac{\text{RMS}[\Delta\theta]_{t-1}}{\text{RMS}[g]_t} g_t$
 Accumulate Updates: $E[\Delta\theta^2]_t = \rho E[\Delta\theta^2]_{t-1} + (1 - \rho)\Delta\theta_t^2$
 Apply Update: $\theta_{t+1} = \theta_t + \Delta\theta_t$
end for

Figure 7: Algorithm of ADADELTA. Note that $\text{RMS}[x]_t = \sqrt{E[x]_t + \epsilon}$ as in (Becker et al., 1988). The hyperparameters ρ and ϵ do not alter the performance of the model significantly; Source: (Zeiler, 2012).




Figure 8: Illustration of 1) underfitting, 2) a good approximation and 3) overfitting; Source: Author's own representation

(2015) offer, therefore, to solve these two issues by doing an initialisation as follow:

$$W \sim U(0; \sqrt{\frac{2}{d_{in}}})$$

With $N(a, b)$ being a normal distribution of mean a and standard deviation b .

Generalisation challenges and regularisation

As defined at the beginning of this section, the accuracy of the classification algorithm is how often the couple $\langle d_j, c_i \rangle$ matches the values of the pre-classified corpus. Also, it was mentioned that the data set is usually split between a training set and a test set. When training the algorithm, we therefore obtain a training error - the proportion of examples for which the model produces an incorrect output. Similarly, we obtain a test error, when running the model on the test set. One challenge in training a model is to avoid a training error that is too high - problem named underfitting - which is the result of a high bias. It produces a model that is too general and not capable of proper predictions with unknown inputs. A second challenge is to have a gap between the training error and test error to be too wide - which is called overfitting - which makes the model to be too specific to the training set and thus not generalizable for new data. Both problems are illustrated in Figure 8 with an analogy to regressions. It must be pointed out that no classification algorithm exists that outperforms

other on all possible data distribution. It is known as the no free lunch theorem (Wolpert, 1996) which is a generalisation of the inter-indexed inconsistency mentioned earlier in this section. Nevertheless, we can find algorithms that perform well on a specific distribution. As expressed in paragraph 2.1, neural networks are particularly capable of approximating any Borel functions. However, it makes them also particularly prone to overfitting. To minimise it, one could get more and better data or regularize the model.

Regularization "is any modification we make to a learning algorithm that is intended to reduce its generalization error, but not its training error" (Goodfellow et al., 2016, Chapter 5). Regularization is a widely researched topic in machine learning but the most common forms of regularization are weight penalties, early stopping, and dropout.

Weight penalties

Weight penalties consist of adding an element to the loss function $L(\hat{c}, c)$ depending on the magnitude of the weights in the matrix W and a hyperparameter γ controlling for the amount of penalty. Two common weight penalties used are called L1 - also called Lasso regression (Tibshirani, 1996) - and L2 regularisation - also called Tikhonov regularisation or ridge regression (Ng, 2004). Let's define a new loss function $L^*(\hat{c}, c)$, below the equations of L1 and L2:




Figure 9: Illustration of overfitting and when the training should stop; Source: Author's own representation

$$L1 : L^*(\hat{c}, c) = L(\hat{c}, c) + \gamma \sum_{w \in W} |w|$$

$$L2 : L^*(\hat{c}, c) = L(\hat{c}, c) + \frac{\gamma}{2} \sum_{w \in W} w^T W$$

Both methods tend to penalise large values in W by shrinking them towards 0, however, in L2 values are squared due to the matrix multiplication and are therefore more penalised. In machine learning literature, L1 appeared first, but L2 has been outperforming L1 in most cases (Ng, 2004).

Early stopping

Early stopping merely consists of stopping the training session before the model starts to learn too much specificity on the training set. This is achieved by stopping when the validation error starts to become greater than for the previous epoch. Indeed, that would mean that the gap between the validation error and the training error is widening and therefore the model starts to become too specific to the training set as shown in Figure 9.

Dropout

Dropout is a method introduced in (Srivastava et al., 2014) that consists of temporarily removing random neurons of the network as shown in Figure 10. A neuron has a probability p of being removed and the authors suggest starting with a value of 0.5 and then adjust if necessary. The rationale behind it is inspired from the role of sex in evolution (Livnat et al., 2010): sexual reproduction generally involves taking half the genes of the male and half of the women ones forcing the genes to “work” together. Similarly, by dropping out neurons from the network, they are obliged to work with randomly selected neurons. It means that a neuron will not overly rely on a specific underlying neuron and learn to adapt

from different inputs, which is the end goal of regularisation. Empirical studies have suggested that dropout is a very effective method of regularisation, in particular with the ReLu activation function (Dahl et al., 2013; Warde-Farley et al., 2013).

2.3. Convolutional Neural Network

For TC tasks, the input of the neural networks is often a sentence or a set of phrases. These have to be encoded in a vector representation (discussion about it in Section 3). This could easily be achieved by considering the sentence as a bag-of-words. However, this method does not take into account the word order. Yet the meaning of a sentence is highly dependent on the word order. CNNs are designed to take into account the context around each word and therefore avoid to consider the input as a bag-of-words. They have been first used in image recognition and then introduced to the NLP community with the work of Collobert et al. (2011) and then showed excellent results even with shallow architecture (Kalchbrenner et al., 2014; Kim, 2014). Since then, CNNs have been continuously used for TC tasks representing the state-of-the-art of text classification techniques (Agrawal and Awekar, 2018; Georgakopoulos et al., 2018; Le et al., 2018; Salinca, 2017; Sundström, 2017). Zhang et al. (2015) developed a similar model to Kim's working at a character-level rather than word level with results varying from a data set to another. Finally, in (Johnson and Zhang, 2017) a deep pyramid CNN model with 15 weight layers was developed. To avoid extravagant computing costs, they decrease the computation time allowed to perform the task in function of the layer depth (from which the pyramid reference comes from). So far this architecture has been the best performing one on several TC tasks.




Figure 10: a) represents a two-layer neural network. b) is the same network with a dropout deactivating two neurons; Source: Author's own representation

The architecture of CNN is similar to the MLP model introduced in section 2.2. The difference lies in the addition of a convolutional layer and a pooling layer represented in Figure 11.

2.4. Convolutional layer

The convolutional layer is present to extract from the input the most salient information - also called feature (more discussion about it in paragraph 3.3) - around a particular window of h words referred as the filter¹¹ in (Kim, 2014). For a filter of size 2 and the sentence “we unlock the potential of the modern workforce”¹², the convolutional layer extracts the features from “we unlock”, then “unlock the”, then “the potential” and so forth. Similarly, a window of size 3 on the same sentence extracts features in “we unlock the”, “unlock the potential”, “the potential of” etc. For each filter, a feature map is created that, from each extraction, stores the different features.

Formally, the layer receives an input vector $s \in R^s$ constructed from a sentence for instance. A dot product is performed between a vector of weights $w \in R^w$ and each w -gram¹³ in s resulting in a new set of features $e = [e_1 \dots e_n]$. The value of n will change depending on the dimension of s and w . If $s \geq w$ then $n = s - w + 1$ (narrow convolution), else $n = s + w - 1$ (wide convolution) with all the $e_i = 0$ for $i > s$.

In the model presented by Kim, a multichannel architecture has been designed - that is a single layer that applies

multiple filters with different sizes on the input and stores the features. Kalchbrenner et al. (2014) later added multiple convolutional layers in their model.

2.4.1. Pooling layer

After the convolutional layer, a set of features is stored for each filter in the filter map. The pooling layer will simply extract the most important feature in each filter map with a function such as $\max(x)$ called 1-max pooling. This is performed to reduce the size of the output, reducing thus the computation power required. In addition, as it reduces the number of parameters θ (w is included in θ), it reduces the risk of overfitting. Kalchbrenner et al. (2014) replaced the 1-max pooling layer by a dynamic k -max pooling layer in charge of extracting the k most important features from the different feature maps. It is called dynamic as the value of k varies in function of the number of the current layer l , the total number of layers L , the dimension of the input s and k_{last} the number of features that are extracted from the last convolution:

$$k_l = \max(k_{last}, \frac{L-l}{L} s)$$

Several methods of pooling exist also using the average or the summation of the features in e , but $\max()$ is the most widely used.

2.5. Recurrent Neural Network

While basic feed-forward neural networks are not able to take into account the word order, we have seen that, by adding a convolutional layer to the architecture, they become

¹¹The filter has the goal to capture the context

¹²Sentence from www.logmeininc.com/

¹³“we unlock” would be a 2-gram in the sentence “we unlock the potential of the modern workforce”. “we unlock the potential” would be a 4-gram.




Figure 11: Illustration of a convolution and a 1-max pooling layer. Each word in the sentence “we unlock the potential of the modern workforce” is represented by a vector of dimension 7. In green, a filter of size 1 is applied, in blue a filter of size 2, in red a filter of size 3 and in yellow a filter of size 4. For each filter, the result is a filter map. For each filter map, a 1-max pooling operation is applied. As 4 different filters were used, the output is of size 4 since only one feature is extracted from the 1-max pooling layer; Source: Author’s own representation

capable of taking into account the context of a word. However, they are not able to take into account the full context as filter sizes are set as hyperparameters. Also, the size of the input vector has to be fixed and therefore during the pre-processing of the data a padding operation - i.e., setting all the input vectors to the same size - must be performed. It is usually done by setting the vector size as big as the longest input in the data set.

Recurrent Neural Network (Elman, 1990), shown in Figure 12, have been particularly suited to work with textual data (Mikolov et al., 2010; Mikolov et al., 2011) because they allow processing variable-length inputs. They do that by being recurrent as they perform the same task for every element of a sequence. The output is then dependent on the previous computation. Compared to MLPs or CNNs, RNNs have an additional component - a hidden state vector - that memorises the previous information. Using the same sentence “we unlock the potential of the modern workforce”, the model first processes the word “we”, then the word “unlock” taking into account the computation performed for “we”. Then, it processes the word “of” taking into account the computation performed for the word “unlock” which was computed using the computation for the word “we”. The algorithm goes one until the end of the sentence. The output includes, therefore, all the computation performed for every single word in the sentence. We understand therefore, why RNNs have first been used for language modelling (Martens, 2011; Mikolov et al., 2010, Mikolov et al., 2010): if the output of the computation is a conditional probability based on the previous words, they can thus predict the next word (Sundermeyer et al., 2014). Similarly, for TC, after each word the condi-

tional probability of the sentence being in a class category is updated until the end of the sentence. The output represents the probability of the whole sentence being classified in a certain category based on all the words in it.

Formally, for an input vector $x = [x_1, \dots, x_{input}] \in R^{input}$, a scalar v_1 is formed by concatenating the vector representing a word in x , and s_0 is the hidden state at iteration 0. Then for i starting at 1¹⁴

$$\begin{aligned} v_i &= concatenate(x_i) = [x_1; \dots; x_{input}] \\ s_i &= f(v_i W + s_{i-1} V) \\ y_i &= g(S_i V) \end{aligned}$$

Where W and V are weights matrices, f an activation function and g another function that results in a probability distribution $y = [y_1, \dots, y_n]$. In their original paper, Mikolov et al. used a sigmoid function for f and a softmax for g . For classification tasks, the intermediate values of y_i are often ignored and only the final one, x_n , is used as it represents the probability of the whole sentence being in a certain class. From the notation above we can observe the recursive nature of the neural network. If we model the hidden state at the 3rd iteration, the equation would be:

$$\begin{aligned} s_3 &= f(v_3 W + s_2 V) \\ &= f(v_3 W + f(v_2 W + s_1 V) V) \\ &= f(v_3 W + f(v_2 W + f(v_1 W + s_0 V) V) V) \end{aligned}$$

¹⁴Please note that I use the notion $[a, \dots, b]$ to define a set and $[a; \dots; b]$ for the concatenation operation




Figure 12: Representation of a Recurrent Neural Network (RNN) with an input vector of dimension 3; Source: Author's own representation

As the process is iterative, we understand how the size of the input can be flexible. Conceptually, the RNN is very similar to an MLP, however, the number of hidden layers is the same as the dimension of the input. Therefore, a layer is “created” for each word that is present in the sentence that we want to classify.

This structure also has some drawbacks. Indeed, due to their recursive nature, RNNs are often difficult to train as they can become very deep neural networks. As a consequence, they often face the problem of vanishing gradient explained in paragraph 2.2.3.4. Also, when it comes to language modelling or classification tasks, sometimes a big gap between relevant information is found. Indeed, it can be that relevant words are at the beginning and the end of the sentence. Therefore, it would be hard for the last iterations to capture the relevance of the first word as it is “drawn” by all the iterations that have been previously performed (Olah, 2017). For these reasons, several improvements in their architecture have been made to tackle classification tasks. The most commonly used are Long Short-Time Memory (LSTM) (Hochreiter and Schmidhuber, 1997) and its variant Gated Recurrent Unit (GRU) (Cho et al., 2014).

2.5.1. Long Short-Time Memory

Simple RNNs introduced the hidden state layer to memorise information. LSTM has an additional variable that tracks the value of the gradients - the memory cell - and three additional layers to monitor and control the memorising of information commonly named input gate, forget gate and output gate. The different gates can be thought as neurons as introduced in paragraph 2.2. In LSTMs, the hidden state layer and the output are the same and therefore $y_i = s_i$.

Conceptually, the memory cell is an object that is going to be updated during the whole iterative process. For each iteration, the memory cell goes through the output gate which gives the final output $y_i = s_i$. How much information comes from the previous iteration that must be forgotten in the memory cell is monitored by the forget gate. Similarly how much of the new information from the current iteration should be added to the memory cell is monitored by the input gate.

Formally, we define $i, f, o, n \in R^{dim_s}$ vectors referred as input gate, forget gate and output gate and new candidate respectively. Also are defined $c = [c_1, \dots, c_{2*dim_s}]$ the memory cell and c_j the memory components. Then:

$$n_i = \tanh(W_n[c_{i-1}, y_{i-1}; x_i]) \quad (1)$$

$$f o_i = \sigma(W_{f_o}[c_{i-1}, y_{i-1}; x_i]) \quad (2)$$

$$i_i = \sigma(W_i[c_{i-1}, y_{i-1}; x_i]) \quad (3)$$

$$c_i = f o_i * c_{i-1} + n_i * i_i \quad (4)$$

$$o_i = \sigma(W_o[c_i, y_{i-1}; x_i]) \quad (5)$$

$$y_i = \tanh(c_i) * o_i$$

First, the new candidate (1) is computed through a $\tanh()$ function which represents the new information coming from the new word x_i . The $\tanh()$ function is applied to make sure that the values are included in the range $[-1 : 1]$. Then, the forget gate (2) and the input gate (2') are computed simultaneously through a sigmoid function σ . This step, thanks to the sigmoid function, tells that value close to 0 must be forgotten and values close to 1 must be saved. From there, the memory components c_j (3) can be computed from the new candidate, the input gate and the forget gate again with a sigmoid function. The output gate (4) is computed from the new memory component c_j . Finally, the output y_i (5) is computed from the dot product between the \tanh of the memory component and the output gate. y_n therefore represents the final probability distribution over the different classification categories.

2.5.2. GRU

The GRU architecture is a simplification of the classic LSTM model but has shown to be competitive for TC tasks (Berger, 2014). Like normal RNNs, GRUs use a hidden state layer but have an update gate and a reset gate.

Mathematically, we define $u, r, n \in R^{dim_s}$, vector referred as update gate, reset gate and new candidate respectively, then:

$$\begin{aligned}
 u_i &= \sigma(W_u[y_{i-1}; x_i]) \\
 r_i &= \sigma(W_r[c_{i-1}, y_{i-1}; x_i]) \\
 n_i &= \tanh(W_n[y_{i-1} \times r_i; x_i]) \\
 y_i &= (1 - u_i) \times y_{i-1} + u_i \times n_i
 \end{aligned}$$

Similar to the LSTM architecture, the gates monitor the quantity of new information that should be added at each iteration. The output is simply an interpolation between the previous iteration - controlled by the update gate - and the new iteration-controlled by the reset gate through the computation of the new candidate.

Even if LSTMs, GRUs, and variants are better suited for language modelling, they have been able to compete against CNNs for TC tasks (Ding et al., 2018; Lee and Dernoncourt, 2016; Liu et al., 2016; Zhou et al., 2016a). Recently, Yu et al. (2018) successfully mimicked skimming, re-reading and skipping techniques performed by humans during TC tasks with an LSTM design. They achieved that by adding a cost function that is minimised during the whole process, providing a better accuracy and higher efficiency than previous approaches. Also, Ma et al. (2018) provide an extension of LSTM that has a separate output gate that incorporates the explicit knowledge such as common sense facts for accomplishing a specific task. The architecture achieved promising results.

2.6. Comparison

We have seen that CNNs are efficient machines in extracting local features around words, but weak at deriving features from sequential treatments because of their rigid structure. On the other hand, RNNs are effective at learning features from sequential correlations, but unable to do it in a parallel way (Zhou et al., 2015b). The two methods seem complementary and in (Yin et al., 2017) the authors point out that which architecture performs better depends on “how important it is to understand the whole sequence”. Indeed, they found that RNNs are not particularly well suited when critical information has to be identified in a sentence to take a classification decision. It includes identifying a particular word to determine the topic or the sentiment of the sentence. They also note that CNNs and GRUs are comparable when sentences are small (<10), but GRU becomes better when the sentences become longer. Finally, according to Baidu Research DeepBench benchmark¹⁵, CNNs are approximately 5x faster to train than RNNs. The iterative nature of RNNs may explain this result.

As it is not clear which one performs better, Zhou et al., 2015b developed a model combining a convolutional layer and an LSTM one. Their model has been able to outperform both CNNs and LSTMs based models. Xiao and Cho (2016)

also developed a hybrid model made out of a recurrent layer (LSTM) and several convolutional layers. However, the input of their model is not working at word level but at character level. Their model has not been able to outperform either simple CNNs or RNNs model on all common classification benchmarks as their results were highly dependent on the data set.

In this section, the classic methodology of solving text classification problems using machine learning has been introduced. Then, the main components of neural networks namely the neurons, the parameters, activation functions and output layer have been described. From there, an explanation of the training procedure of neural network by initialising the parameters and using a loss function to minimise through a training algorithm has been provided. Finally, regularisation techniques to improve the generalisation power of the neural networks were presented.

Building on the previous explanations, the functioning of CNNs as powerful tools to learn local features thanks to a convolutional layer and a pooling layer has been highlighted. Also, the ability of RNNs to learn sequential features has been explained, and a comparison of both models has been provided together with their most up-to-date applications.

The next section is dedicated to explaining how to convert textual information in a format suitable to be fed in the neural networks described.

3. Document Representation

“Translation is not original creation - that is what one must remember. In translation, some loss is inevitable” Joseph Brodsky

As computers work with binary information, they are not able to directly interpret a human language. Consequently, the second challenge of TC is to determine the best representation of the input for the classifier to extract the syntactic structure and semantics of texts. Indeed, the effectiveness of most classifiers is heavily dependent on the choice of the representation of the data (Bengio et al., 2013; Wolfram and Zhang, 2008). This task is often referred as learning representation (or document indexing).

Several approaches have been developed and are based on the idea that a document can be described based on a set of the words contained in it commonly called the set-of-words or the bag-of-words approach (Apté et al., 1994; Fuhr et al., 1991; Lewis, 1992; Tzeras and Hartmann, 1993). Furthermore, a word has been proved to be the best unit for text representation (Song et al., 2005) despite promising recent results of representations built at character level (Conneau et al., 2016). However, not all words have the same representative value. Indeed, words such as “and” or “or” would not provide the same information as “music” or “image” about the topic of a document. A solution has therefore been to develop a vector representation of the document where the “importance” of each word is stored. Determining such importance has been a highly investigated field of NLP and will

¹⁵<https://github.com/baidu-research/DeepBench#results> The website compares different hardware components for data science tasks including training RNNs and CNNs.

be discussed in 3.1. In the past decade, various methods have been approached and the currently predominant one is the vector space model (VSM) introduced by Salton et al. (1975).

3.1. The vector space model

In the vector space model, documents are represented as a vector where each dimension represents a separate term (i.e., word), and weights are ranging between [0, 1]. 0 is used to express the absence of a term in the document and all value bigger aim to represent the importance of the word in the document.

For $D = \{d_1, \dots, d_n\}$ a set of documents, we define $L = \{l_1, \dots, l_m\}$ being the dictionary (or lexicon), i.e., the set of all different terms occurring in D. Then we define, a document vector as $d_i = \langle w_{1i}, \dots, w_{ni} \rangle$ with w_{ki} representing the weight of the k^{th} term in d_i . Given the vector documents for two documents, it is then possible to determine the similarity - product of vector or inverse function of the angle between the two vectors - between them (Salton et al., 1975). Also, to give all the documents the same importance, each vector document is normalized to have lengths of one.

Encoding the vectors, i.e., determine the weight w_i of a word l_i in a document d_j has been subjects to many discussions (Baeza-Yates and Ribeiro-Neto, 1999; Gövert et al., 1999), but a common approach has been to use the tfidf function introduced in (Salton & Buckley, 1988):

$$w_i(d_j, l_i) = \frac{tf(d_i, l_i) \log\left(\frac{N}{n_t}\right)}{\sqrt{\sum_{j=1}^m tf(d_i, l_i)^2 \left(\log\left(\frac{N}{n_t}\right)\right)^2}}$$

With N being the number of documents in D, n_t the number of documents in D that have an occurrence of l and $tf(d, l)$ the number of time l appears in d. With such a method, deriving the similarity between two documents d_1 and d_2 becomes handy as it can be represented by the Euclidian distance between the two document vectors d_1 and d_2 .

However, the drawback is the high dimensionality of the representation. Indeed, for a set D of size N with M unique words in L, the matrix representation is of size NxM whose rows are words and columns are documents (Sánchez et al., 2008). To overcome this issue, some pre-processing can be done on the data which is discussed in the next paragraph.

3.2. Tokenization, filtering and stemming

As exposed before, the most common unit in text classification task is the word. Therefore for each document d, a tokenization is required, i.e removing all punctuations marks and replacing non-text characters by single white spaces (Murty et al., 2011). It has been highlighted that by representing the set of documents on a VSM, we end up with a representation that has a high dimensionality. To reduce it, the first method is to diminish the size of the lexicon L. This can be done by filtering, i.e., removing words from the

lexicon. Frakes (1992) point out that words that appear really often bear no particular statistical relevance and can be removed. Also, words such as prepositions or articles do not have content information. In addition, stemming can be performed on the data which consists of grouping words with the same roots and replacing it with the most basic form or stem. It is indeed assumed that words with a common stem will usually have similar meanings (Porter, 1980). Therefore plural forms from nouns or the “ing” from verbs will be removed and the dictionary will contain a list of unique stems.

3.3. Distributed representation of words

Although pre-processing techniques have been able to reduce the dimensionality of the document representation efficiently, the modelling presented earlier has other drawbacks. They include not being able to represent the distance between individuals terms (Kusner et al., 2015) that means it does not capture sense about the semantics of the words. Also, the high dimensionality is often not suitable for computing document distance as they produce matrixes that are almost orthogonal (also called diagonal dominance¹⁶) (Greene and Cunningham, 2006). Finally, word order is disregarded when constructing such a representation. Some studies have been trying to solve this issue producing a more coherent approach, yet without improving the performance of the downstream classification task. (Blei et al., 2003; Deerwester et al., 1990; Robertson and Walker, 1994).

A breakthrough in document representation occurred when researchers leveraged the distributional hypothesis that states that words that are used and occur in the same contexts tend to purport similar meanings (Harris, 1981). Additionally, the pioneering work of Hinton et al. (1986) on distributed representations contributed to improvement of document representation: rather than representing a word with a single high dimensionality vector, it can be represented as a combination of low dimensional vectors. Each vector is used to represent a feature (such as the tfidf of a word, Chi-Squared, Information Gain (Debole and Sebastiani, 2003)) and the number of features is smaller than the size of the lexicon (reducing thus the dimensionality). Also, relevant features can be selected from a set of all the features (feature selection) and used for the representation. Alternatively, a machine learning approach can be implemented to pick and transform the features (feature extraction) into a lower dimension. This distributed representation of a word is called word embedding. A word is thus represented as a word vector with each dimension representing a feature.

Formally, for each word l in L, a set of linguistic features $[e_1, \dots, e_k]$ is extracted or constructed. Each e_i is encoded in a vector $v(e_i)$. l is then represented by a combination of each vector (summation, concatenation or both). The model is therefore made out of dense and low-dimensional vectors

¹⁶A square matrix A is called diagonally dominant if $|A_{ii}| \geq \sum_{i \neq j} |A_{ij}|$ for all i.




Figure 13: Feature filter model; Source: (John et al., 1994)




Figure 14: Feature wrapper model; Source: (John et al., 1994)

which lowers the dimension of the representation of the document significantly. These vectors are usually the input of the classifier mentioned in paragraph 2.2.1.

3.4. Feature selection

In the previous section, the notion of feature for words has been introduced. A plethora of features exists concerning words and to generate the representation of a document, some criteria can be used to filter out if a feature is relevant¹⁷ for prediction purpose or not. Feature selection methods are categorized in three different types: filter, wrapper, and embedded methods.

3.4.1. Filter

Filter method (illustration of process in Figure 13) refers to algorithms that treat a possible set of features and rank them independently of the classifier. The top-ranked features are selected (Forman, 2003). Examples of such algorithm include some built on similarity measures such as Pearson's correlation coefficient (Saeys et al., 2008), statistical methods and heuristic search or ensemble learning (Kira and Rendell, 1992). These methods have the advantage of being fast and thus scalable. However, they do not yield particularly accurate results as they increase bias and are exposed to the selection of redundant features (Jashki et al., 2009).

3.4.2. Wrapper

Wrapper methods (illustration of the process in Figure 14), on the other hand, test every feature in the context of the classifier (Kohavi and John, 1997). They usually involve automated search techniques such as the greedy search strategy (Guyon and Elisseeff, 2003). These methods are more accurate than filter methods but come with high-computing costs.

¹⁷Discussions about the meaning of relevance and its definition can be found in (Sag et al., 2002)

3.4.3. Embedded

Finally, embedded methods perform feature selection during the execution of the classifier (being therefore embedded in the classifier). Therefore, the feature selection and the training methods of the classifier are not separated steps. Conventional methods may use decision tree algorithm (Genuer et al., 2010) or multinomial logistic regression (Cawley et al., 2007). These methods are similar to wrappers but are specific to classifiers, which makes them computationally less expensive as they are optimized for them.

3.5. Feature extraction

As mentioned, a board range of features exists and some that humans find useful will not necessarily be useful for the models and vice-versa. Therefore, all the features known based on basic statistics about a document can be used, referred as count based methods. Alternatively, machine learning techniques such as neural networks can be used to let the model determine which features are important or not.

3.5.1. Count based methods

In feature extraction, the feature space - set of all possible features - is converted to another space with a lower dimension keeping the most informative and discriminative features (Gomez et al., 2012). Methods include Principal Components Analysis (PCA) and Latent Semantic Analysis (LSA).

PCA is a statistical method that transforms the set of features (possibly correlated) into new features that are uncorrelated called principal components using a linear transformation. Like in feature selection, the best new features are then selected.

LSA (Deerwester et al., 1990) - also referred as Latent Semantic Indexing - is a technique developed to address the problems deriving from the use of synonymous, near-synonymous, and polysemous words as features of document representations (Sebastiani, 2002). The process involves

identifying the relevant words - using, for example, the tfidf of words - and then constructs a term-document matrix as described in paragraph 3.1. Then the matrix is decomposed using Singular Value Decomposition - a technique closely related to PCA. The result is a set of lower dimension features vectors that were constructed looking at patterns of word usage in the documents. In essence, the features are usually hardly interpretable as there are meant to capture latent (hidden) relationship between words. LSA provided a significant step forward in document representation as it accounted for semantic characteristics of texts, synonymy of words and partially polysemy (Deerwester et al., 1990).

Glove: the state-of-the-art of count-based model

In the paper introduced by Pennington et al. (2014), the authors argue that the count of words in a document carries meaningful information, but also the count of a word w_i in the context of another word w_j called co-occurrence probability. Following their example, for the context of steam and ice, it is expected that the ratio of the probability of observing solid in the context of ice and the probability of observing solid in the context of steam - $\frac{p(\text{solid}|\text{ice})}{p(\text{solid}|\text{steam})}$ - to be high. Likewise, this ratio for the word gas in the same contexts is expected to be small. The model therefore constructs a matrix X_{ij} based on word-context co-occurrences and factorise it to obtain the vectors. To complete the latter step, the authors use a weighted least squares regression model that is able to encode the information available in the probability of co-occurrence. When constructing the word vectors, the objective is to minimize the difference between the product of the two word vectors w_i and w_j (word and context), and the logarithm of the probability of co-occurrence (plus a bias for each word) which is expressed as follow:

$$J = \sum_{i,j=1}^V f(X_{i,j})(w_i^T w_j + b_i + b_j - \log(X_{i,j}))^2$$

$$\text{with } f(X_{i,j}) = \begin{cases} (\frac{X_{i,j}}{X_{max}})^{3/4} & \text{if } X_{i,j} < X_{max} \\ 1 & \text{otherwise} \end{cases}$$

3.5.2. Neural networks for words embedding

Methods to represent words explained so far are referred as count-based methods in (Baroni et al., 2014) as values in vectors are derived from co-occurrence counts. The authors point out the weaknesses of these models namely problem of scalability, poor performance on word analogy evaluation and task-dependent (except for GloVe that performed pretty well on the latter). To deal with these issues, new models have appeared referred as predictive-based methods. This new generation of models were first exposed in 1981 (Hinton et al., 1986), but have demonstrated their utility in (Collobert and Weston, 2008) building up on previous research on deep neural network (Bengio et al., 2003) challenging the previous state-of-the-art methods. Rather than counting words co-occurrence, generating the vectors and reduc-

ing the dimensionality, these methods try to directly generate the vectors by predicting a word from its neighbours or vice versa. Thus, as similar words occur in similar contexts, the system assigns similar vectors to similar words (Baroni et al., 2014). The comparisons between count-based and predictive-based methods have demonstrated the superiority of the latter in lexical semantics tasks including semantic relatedness, synonym detection and analogy, (Cambria et al., 2017; Socher et al., 2011; Turney and Pantel, 2010; Weston et al., 2010), but have failed to leverage statistical information from documents as they are based on context windows of a few words. It must however be pointed out that no methods of adequately evaluating the quality of vector representations have been developed. Indeed, so far they have been evaluated on word similarity or analogy metrics, but these only correlate weakly with downstream tasks performance such as TC (Tsvetkov et al., 2015).

The next section is dedicated to presenting the most famous predictive-based methods using neural networks.

Word2Vec

In the paper (Mikolov et al., 2013), the authors offer two models. One model predicts a word given a context (Continuous Bag-of-Words model) and the other one given a context, predicts a word (Skip-gram model). Using these models with such objectives will not result in word vectors per se in the output layer. Indeed, the word vectors will be present in the different weight matrices of the models. The intuition behind it was previously expressed: if two words are similar, they should appear in a similar context and thus their representation should be similar.

The results of the learned embedding were a big step forward in the vector representation of words. Indeed, they were not only capable of training a huge list of words (1.6 billion) in less than one day, but also captured semantic meaning of words. It is illustrated by an example that has become famous in the NLP community: having the vector for the words queen, women, men and king, they have performed the following calculation successfully:

$$\text{queen} - \text{women} + \text{men} = \text{king}$$

In the same vein, they were capable of capturing the semantic behind the sentence "France is to Paris as Germany is to Berlin".

Continuous Bag-of-Words model

To achieve this amazing result, they leverage a modification of the MLP presented in 2.2 as pictured in Figure 15. They used the same structure with an input layer that represents one-hot-encoded¹⁸ words, a single hidden layer and an output layer with the goal to classify. They looked at four words and given these words try to find a word that would fit in the middle of these four words which can be defined as a classification problem. With this architecture, they therefore end

¹⁸A one-hot-encoded vector consists of 0s in each dimension with the exception of a 1 in a dimension used uniquely to identify the word by its position in the sentence.




Figure 15: x_1, x_2, x_3 is the input layer made out of one-hot-encoded vectors. The hidden layer is represented by neurons and the output layer is y_1, y_2, y_3 . This representation would fit a sentence made out of three words. Source: Author's own representation

up with two weights matrices W^1 (from the input layer to hidden layer) and W^2 (from the hidden layer to the output layer) that form the parameters θ . The output layer, thanks to a softmax function, represents a multinomial distribution of all words given a context¹⁹. We understand that the goal is to maximize the probability of a word given a context, which consists of minimizing the opposite probability:

$$\text{Lossfunction} = -\log(p(\text{word}|\text{words}_{\text{context}}))$$

The log appears because we are using the softmax function to transform the last layer in a probability distribution.

With this setting, the actual output of interest are the matrices W^1 and W^2 . Indeed, after training, the matrix W^1 contains in its lines vectors that, for a word, represents the context. On the other hand, W^2 has a vector representation of a word in its columns, which is precisely what we are looking for (Rong, 2014).

Skip-Gram model

The Skip-Gram model is very similar to the CBOW model. It is just doing the opposite: given a word, predict the context. Indeed, for a word given, it will pick another word and estimate the probability of that word being around²⁰ it. Consequently, the rows of W^1 will now represent the vector representation for a word and the column of W^2 will represent context vectors.

From the two models, the authors have been able to create vectors that were capable of representing words better

syntactically (with the CBOW model) and semantically (with the Skip-gram model) than previous neural models (Mikolov et al., 2009; Mikolov et al., 2010). However, the models have limitations. The first one is that for one word, they assign one vector and therefore they are unable to represent polysemy words. To (partially) solve the issue, Upadhyay et al. (2017) developed an algorithm that learns word representation jointly across language. The intuition behind it is that a polysemy word in language could be translated into distinctive words in another language. Using the authors' example, the word bank in English which has several meanings can be translated to banque or banc in French which capture two different meanings with two different words. Therefore, by learning using multiple languages, the algorithm can identify which sense to use. The second caveat is that the meaning of multi-word expressions²¹ is not captured. Indeed, expressions such as "in short" or "Los Angeles" are poorly encoded as they will be represented in two vectors. Some methods have been developed to capture phrasemes without however improving the performance of downstream tasks such as text classifications (Hashimoto and Tsuruoka, 2016; Yu and Dredze, 2015). Finally, training the CBOW or Skip-gram is computationally expensive on large datasets. Thus, rather than generating an embedding for every task performed, it is common practice to use pre-trained vectors²². However, it is often the case that some words in the datasets are not part of the pre-trained vectors. These words are referred in the literature as out-of-vocabulary words (OOV). Being able to assign a proper representation to the input, including OOV

¹⁹The context can be made out of one word or several words preceding or following the word of interest

²⁰"around" is predefined and can be for instance 2 words before and 2 words after. The authors of the model found that increasing the size of the context resulted in better quality of word vectors.

²¹Multword expressions are also called phraseme. An accurate discussion about the typology of multiword expression is present in (Sag et al., 2002)

²²Mikolov et al., after developing their model, published a set of pre-trained vectors on Google News with 3 million words and phrases.

words, can alter the performance of the downstream task by up to 6% over random initialization (Dhingra et al., 2017). One way of dealing with OOV words is to replace them with a unique token, UNK (Chen et al., 2016; Shen et al., 2017), and use it for training. Another method is to assign each OOV word a randomly generated vector at test time (Kim, 2014) or a unique randomly generated vector (Dhingra et al., 2017). A recently suggested method in (Dong and Huang, 2018) is to combine pre-trained vectors and vectors generated during training. When a word is present in the pre-trained vectors and the training set, then a new vector is constructed concatenating both vectors. If one of them is missing, it is replaced by a null vector.

Proposal

I would like to propose a variation of the method developed in (Dong and Huang, 2018). Rather than pre-trained or embedded vectors concatenating with a null one, I suggest to concatenate them with a unique vector sampled from a distribution such as the vector has the same variance as the other ones. The idea of generating vectors from such a distribution is not new as it was already expressed in (Kim, 2014). However, I combine both approaches with a concatenation operation as shown in Figure 16.

FastText

As expressed earlier, training word vectors can be computationally expensive to learn and dealing with OOV words can be challenging. Bojanowski et al. (2016) offer an extension of Word2Vec named FastText to learn a vector representation of word quickly and to (partially) deal with OOV words. Rather than learning vector representation of words, they learn representations of character n-grams. Then a word is simply the sum of this character n-gram²³. For instance, for the word “hello”, extracting a character 3-gram, will give the vector representations of: “he”, “hel”, “ell”, “llo”, “lo”. This allows leveraging the morphology of words and therefore reducing the number of necessary computations. Also, when dealing with OOV words, it is likely that new words can be expressed as a combination of the learned character n-gram.

While Section 1 described various neural networks models, Section 2 of this work has been first dedicated to explaining the rationale behind the conversion of words into vectors as inputs for the models. From the simple bag-of-words method that uses a high dimensional representation, the notion of feature and tricks to diminish the size of that representation have been introduced. Furthermore, methods to select features but also techniques that extract them were explained. For the latter, the dichotomy that exists between count-based solutions - with its best representative GloVe - and prediction-based solutions such as Word2Vec and FastText has been presented. Finally, a solution to deal with words that are not present in pre-trained vectors data set has been proposed.

As the state-of-the-art of neural networks models for TC and word embedding methods have been identified, the next

section describes the benchmark used to evaluate them on the LogMeIn data. Also I introduce another dataset to assess the proposal.

4. Experiment

“Experience is simply the name we give our mistakes” Oscar Wilde

To evaluate the state-of-the-art classifiers on the data provided by LogMeIn, the CNN model as described in (Kim, 2014) and a hybrid CNN+LSTM model as described in (Zhou et al., 2015b) were implemented. Also, in order to evaluate the proposal, the implementation was tested on two datasets: the LogMeIn one and the TREC (Li and Roth, 2006) dataset.

Moreover, I test different techniques of embedding, a random initialization with different dimensions, using pre-trained vectors generated by Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), FastText (Bojanowski et al., 2016), FastText with subwords information (Mikolov et al., 2017), the method introduced in (Dong and Huang, 2018) and my proposal.

4.1. Data

The LogMeIn dataset is made out of customer reviews based on the product GoToMeeting - an online meeting and video conferencing software. It has been annotated such as reviews are classified under the categories “screen”, “video”, or “audio”. Unfortunately, reviews may appear in several categories in the original dataset. Therefore the dataset is used for a binary classification problem whether the review is under the category “audio” or not. Also, to avoid bias, an under-sampling procedure has been performed on the “non-audio” category to get a 1:1 ratio of “audio” and “non-audio” entries. It consisted of randomly dropping data points until parity was reached.

The TRAC dataset is a common benchmark used for multi-topic categorisation. It is made out of a question that refers to a person (884 samples), a location (616), numeric information (664), an abbreviation (62), and an entity (937). The task is to classify a question under one of these categories. Statistics for both datasets are available in Table 3.

4.2. Models

The CNN is the same as described in Section 2.3. As explained, a CNN takes a fixed length of input; therefore each sentence is padded to the maximum sample length. It is done by adding symbols for sentences that are shorter than the maximum sentence length in the training set and taking off words for samples that are longer in the test set. The window sizes of the filters of the convolutional layer are 3, 4 and 5 words with 100 feature maps each which are combined with a ReLU non-linearity. A 1-max pooling operation is performed on the output of the filters and then a 0.5 dropout is applied. The batch size is 32 and the loss function is the softmax cross entropy function shown below:

²³In their study, they extract all the n-gram with n ranging from 3 to 6.

Algorithm : Combine pre-trained word embedding with those generated on training set.

Input : Pre-trained word embedding set $\{U_w|w \in S\}$ where $U_w \in \mathbb{R}^{d_1}$ is embedding vector for word w . Word embedding $\{V_w|w \in T\}$ are generated on training set where $V_w \in \mathbb{R}^{d_2}$. P is a set of word vocabulary on the task dataset. a such as $U[a; a]$ has the same variance as the vectors in $\{V_w|w \in T\}$. b such as $U[b; b]$ has the same variance as the vectors in $\{U_w|w \in S\}$

Output: A dictionary with word embedding vectors of dimension $d_1 + d_2$ for $(S \cap P) \cup T$.

```

res = dict()
for w in (S ∩ P) ∪ T do
  if w in S ∩ P and w in T then res[w] = [U_w; V_w];
  else if w in S ∩ P and w not in T then V_w ~ U[a; a] and res[w] = [U_w; V_w];
  else U_w ~ U[b; b] and res[w] = [U_w; V_w];
end
Return res

```

Figure 16: Algorithm suggested to deal with OVV word, using the same notation as in (Dong and Huang, 2018); Source: Author's own representation

Table 3: Statistics about LogMeIn and TREC datasets Source: Data compiled by author

	LogMeIn	TREC
Total Samples	1137	4000
Train Samples	1023	3600
Test Samples	114	400
Number of categories	2	6
Average length of samples	16.975	10.1545
Median length of samples	10	10
Maximum length of samples	205	37
Total unique words	2786	6987

$$L(\hat{c}, c) = -\log\left(\frac{e^c}{\sum_{i=1}^C e^{\hat{c}_i}}\right)$$

The CNN+LSTM model is first made out of convolutional layer that extracts higher-level sequences of word features. It is the same convolutional layer as the simple CNN model. Unlike in (Zhou et al., 2015a) a 1-max pooling operation is kept after the convolutional layer. Then an LSTM capture long-term dependencies over each window feature created by the convolutional layer. After the LSTM, a 0.5 dropout is applied just before the softmax cross entropy layer. The batch size is also 32.

4.3. Word embedding

I test 8 forms of word embedding. First, I try two random assignments of vectors to words from the uniform distribution $U[-0.25, 0.25]$. The first set of vectors is of dimension 300 and the second of dimension 600. This is to check the effect of the size of word embedding on the downstream classification task.

Also, a third embedding is generated from pre-trained vectors with Word2Vec made available by Mikolov et al. (2013). It includes a vocabulary of 3 million words and phrases that were trained on about 100 billion words from a Google News dataset. The vectors are of dimension 300.

The fourth embedding is generated from pre-trained vectors with GloVe made available by Pennington et al. (2014). They were trained on a Wikipedia and Giga word datasets²⁴ and consist of 400'000 words. The vectors are of dimension 300.

The fifth and sixth embedding are generated from pre-trained vectors with FastText (Bojanowski et al., 2016; Mikolov et al., 2017). They consist of 1 million word vectors trained on Wikipedia 2017²⁵, UMBC web base corpus²⁶ and statmt.org news dataset²⁷. One is trained with subword information the other is not.

The sixth embedding method is the proposal of Dong and Huang (2018) which consists of vectors of dimension 600 made out of the concatenation of Word2vec pre-trained vectors and vectors trained directly on the database with the CBOW algorithm. If a word is not present in either of the two sources it is represented by a null vector of size 300. Finally, the last embedding is my proposal. It is the same as the sixth except that vectors that are not found are represented by a unique vector initialized from a uniform distribution $U[-0.25, 0.25]$ of size 300.

The code was written in Python using TensorFlow²⁸ and

²⁴<https://catalog.ldc.upenn.edu/LDC2012T21>

²⁵<http://wiki.dbpedia.org/Datasets>

²⁶<https://ebiquity.umbc.edu/resource/html/id/351>

²⁷<http://www.statmt.org/>

²⁸<https://github.com/tensorflow/tensorflow/releases/tag/v1.8.0>

the code of Jie Zhang available on GitHub²⁹ as a basis for the implementation of the CNN and CNN+LSTM. Also, the gensim³⁰ implementation of Word2Vec is used to generate the word embedding of the TREC and LogMeIn dataset. The code is available in Appendix 8.1, Appendix 8.2, and Appendix 8.3. Finally, the tests were performed on a Central Processor Unit (CPU) Intel (R) Core™ i7-7500U @ 2.70 GHz and 8GB LPDDR3-1866Mhz RAM. As some randomness is part of each model, they are tested 5 times. The average performance as well as a 95% confidence interval is reported. The 16 models, their name, and characteristics are summarized in Table 4.

5. Results and Discussion

“I am just a child who has never grown up. I still keep asking these ‘how’ and ‘why’ questions. Occasionally, I find an answer” Stephen Hawking

The first part of this section presents the results necessary to compare the effectiveness of the CNN and the CNN+LSTM models. Then, the effects of the different word embedding methods are presented and discussed.

5.1. CNN and CNN+LSTM

A first remark is that I was not able to achieve the same results as (Kim, 2014) and (Zhou et al., 2015a) on the TREC dataset. As pointed out in Wang et al., (2018), measures can change depending on the pre-processing of the data. In my experiment, short forms such as “I’m” or “He’ll” are split in two distinctive words, uppercase characters are replaced by lowercase ones and non-alphanumeric characters except punctuation symbols were removed. Also, I do not use the same hyperparameters and architecture. Indeed, when Kim use ADADELTA (Zeiler, 2012) as the algorithm to update gradients, I use ADAM (Kingma and Ba, 2015). Indeed, it was shown in the aforementioned paper that both methods efficiently lower the cost of training on CNNs, but ADAM is better at that task than ADADELTA, especially on deep neural networks. However, despite being more efficient, ADAM and ADADELTA should converge toward the same local minimum which should therefore not change the performance of the downstream task. Tests with an ADADELTA function have been performed on the TREC dataset and the CNNFXT model and no significant changes were perceived confirming the previous statement. Also, Kim uses 25 training cycle (or epochs) whereas on my benchmark I only use 3. The more epoch is used, the better trained the model is, however the higher the risk of overfitting. I have personally chosen 3 epochs to run more tests in the benchmark. Indeed, multiplying the amount of epochs inevitably increases the training time of each model. However, after testing on the CNNW2V

model on the TREC dataset with 25 epochs, again no significant differences were observed as it appears that the model plateaus at about 84% accuracy. Despite these differences, I have not been able to identify other sources responsible for the performances differences.

Similarly, the results of the CNN+LSTM models do not replicate the ones from C. Zhou et al. The first difference is the presence of the max pooling layer after the convolutional layer. The authors argue that the operation breaks the sequence order as the selected features from the convolutional layer are discontinuous. However, the role of the max-pool is first reducing the computation for the next layer, but also to extract the most salient features in the sample. A test has been performed on the LSTMW2V model, but again no significant changes in term of performance have been observed.

The number of epochs, however, affects much more the LSTM models. Indeed, the results reported in Table 5 come from a training procedure of 3 epochs, but by increasing it to 25, the accuracy for the model LSTMW2V jumped from 63% and 77% to 86% and 97% on TREC and LogMeIn datasets respectively³¹. To investigate it, I changed the gradient updating algorithm of the CNN+LSTM models. In their initial configuration (and the one tested in this work), the algorithm used for the update of the gradient of the CNN+LSTM models is the RMSprop (Tieleman and Hinton, 2012). A test has been performed using the ADAM algorithm with 3 epochs on both datasets and results are conclusive achieving similar results than with the RMSprop with 25 epochs. Therefore, the first recommendation when using a CNN+LSTM model is to use the ADAM algorithm as gradient update function. It requires less training time while yielding better results than RMSprop for the same number of epochs.

In the lights of the first conclusion, a second benchmark has been performed using a CNN+LSTM with an ADAM gradient update function and 3 epochs to compare the model directly with the simple CNN models. The results are reported in Table 6.

Besides, for both configurations, I use filter sizes of 3, 4 and 5 on the convolutional layer. In their original paper, C. Zhou et al. conclude that a filter of size 3 yields better results for the CNN+LSTM architecture. However, Kim reports better results using filter size of 3, 4, 5 on a simple CNN one. I find better results using filter size of 3, 4 and 5 on both datasets with both configurations. As few data were collected (5 per model), to investigate differences between CNN and CNN+LSTM, I aggregate the measures of all CNN tests and all CNN+LSTM tests. The mean and a 95% confidence interval are reported in Table 7.

There are no statistical differences observed between the two models on the dataset experimented. The CNN results are similar in magnitude to the results in (Zhou et al., 2016b) on the TRAC dataset, but no improvement is observed by adding the LSTM layer to the architecture unlike in (Zhou et al., 2015a). First, a better fine-tuning of the CNN+LSTM

²⁹<https://github.com/jiegzhan>

³⁰<https://radimrehurek.com/gensim/>

³¹These figures might be inflated as I did not check whether an overfitting problem was appearing or not.

Table 4: Summary of the different model tested and their features; Source: Data compiled by author

Name	Model	Dimension	Embedding
CNN300	CNN	300	Random
CNN600	CNN	600	Random
CNNW2V	CNN	300	Pre-trained Word2Vec (Mikolov et al., 2013)
CNNNGVE	CNN	300	Pre-trained GloVe (Pennington et al., 2014)
CNNFXT	CNN	300	Pre-trained FastText (Bojanowski et al., 2016)
CNNFXT_SUB	CNN	300	Pre-trained FastText (Mikolov et al., 2017)
CNNW2V600_NULL	CNN	600	Word2Vec + pre-training on dataset (Dong and Huang, 2018)
CNNW2V600	CNN	600	Pre-trained Word2Vec + pre-training on dataset (proposal)
LSTM300	CNN+LSTM	300	Random
LSTM600	CNN+LSTM	600	Random
LSTMW2V	CNN+LSTM	300	Pre-trained Word2Vec (Mikolov et al., 2013)
LSTMNGVE	CNN+LSTM	300	Pre-trained GloVe (Pennington et al., 2014)
LSTMFXT	CNN+LSTM	300	Pre-trained FastText (Bojanowski et al., 2016)
LSTMFXT_SUB	CNN+LSTM	300	Pre-trained FastText (Mikolov et al., 2017)
LSTMW2V600_NULL	CNN+LSTM	600	Pre-trained Word2Vec + pre-training on dataset (Dong and Huang, 2018)
LSTMW2V600	CNN+LSTM	600	Pre-trained Word2Vec + pre-training on dataset (proposal)

Table 5: Classification accuracy of the different models on the LogMeIn and TREC datasets. The best result is in bold; the second best is in italic. The 95% confidence interval is reported in parentheses. Here the CNN+LSTM models are trained with RMSprop; Source: Data compiled by author

	LogMeIn	TREC
CNN300	0.9035 (\pm 0.0241)	0.7765 (\pm 0.0169)
CNN600	0.8614 (\pm 0.0190)	0.7810 (\pm 0.0107)
CNNW2V	0.9333 (\pm 0.0253)	0.8425 (\pm 0.0831)
CNNNGVE	0.9123 (\pm 0.0108)	0.7730 (\pm 0.0162)
CNNFXT	0.9386 (\pm 0.0139)	0.8455 (\pm 0.0161)
CNNFXT_SUB	0.9193 (\pm 0.0385)	0.8300 (\pm 0.0121)
CNNW2V600_NULL	0.9351 (\pm 0.0268)	0.8445 (\pm 0.0176)
CNNW2V600	0.9273 (\pm 0.0268)	0.8165 (\pm 0.0099)
LSTM300	0.6069 (\pm 0.0478)	0.7000 (\pm 0.0438)
LSTM600	0.5825 (\pm 0.0321)	0.7005 (\pm 0.0231)
LSTMW2V	0.6316 (\pm 0.0311)	0.7745 (\pm 0.0348)
LSTMNGVE	0.7175 (\pm 0.0419)	0.7380 (\pm 0.0185)
LSTMFXT	0.6070 (\pm 0.0575)	0.7550 (\pm 0.0360)
LSTMFXT_SUB	0.5316 (\pm 0.0476)	0.6835 (\pm 0.0176)
LSTMW2V600_NULL	0.6386 (\pm 0.0228)	0.7845 (\pm 0.0127)
LSTMW2V600	0.5912 (\pm 0.0275)	0.7535 (\pm 0.0261)

model is necessary. As pointed out in the previous paragraph, LSTM based model are very sensitive to the number of epochs and update algorithm function. Further investigations must be performed on the CNN+LSTM model to identify the right number of epochs, but also the ideal batch size. Indeed, a test has been conducted with the LSTMW2V model with 6 epochs showing an accuracy of 84.25% on TREC, which is higher than all other tests (Appendix 8.4).

In both models, the convolutional layer performs the same task which explains the similarity of results, but the addition of the LSTM layer requires further work to leverage the memory cell capacity. Also, as pointed out in (Yin et al., 2017), CNNs and RNNs are expected to yield comparable results when sentences are short which is the case as shown

in Table 3. Finally, both algorithms perform better on LogMeIn than TREC, but tasks are also slightly different as one is a binary classification and the other one is a 6 categories classification task.

5.2. Effect of Word Embedding

Despite not being able to replicate other state-of-the-art results, effects regarding the word embedding are captured by both models by holding the rest of the parameters constant. To capture only the effect of the word embedding method, an aggregation has been made between results of CNN and CNN+LSTM based models. Results on the LogMeIn dataset is present in Figure 17 and results on TREC are available in Figure 18.

Table 6: Classification accuracy of the different models on the LogMeIn and TREC datasets. The best result is in bold; the second best is in italic. The 95% confidence interval is reported in parentheses. Here the CNN+LSTM models are trained with ADAM; Source: Data compiled by author

	LogMeIn	TREC
CNN300	0.9035 (\pm 0.0241)	0.7765 (\pm 0.0169)
CNN600	0.8614 (\pm 0.0190)	0.7810 (\pm 0.0107)
CNNW2V	0.9333 (\pm 0.0253)	0.8425 (\pm 0.0831)
CNNGVE	0.9123 (\pm 0.0108)	0.7730 (\pm 0.0162)
CNNFXT	0.9386 (\pm 0.0139)	0.8455 (\pm 0.0161)
CNNFXT_SUB	0.9193 (\pm 0.0385)	0.8300 (\pm 0.0121)
CNNW2V600_NULL	0.9351 (\pm 0.0268)	0.8445 (\pm 0.0176)
CNNW2V600	0.9273 (\pm 0.0268)	0.8165 (\pm 0.0099)
LSTM300	0.87016 (\pm 0.0331)	0.7855 (\pm 0.0201)
LSTM600	0.89474 (\pm 0.0300)	0.7895 (\pm 0.0082)
LSTMW2V	0.8895 (\pm 0.0083)	0.835 (\pm 0.0271)
LSTMGVE	0.91924 (\pm 0.0162)	0.816 (\pm 0.0141)
LSTMFXT	0.90596 (\pm 0.0162)	0.8365 (\pm 0.0180)
LSTMFXT_SUB	0.91756 (\pm 0.0122)	0.8145 (\pm 0.233)
LSTMW2V600_NULL	0.91754 (\pm 0.0176)	0.825 (\pm 0.0258)
LSTMW2V600	0.91404 (\pm 0.0100)	0.833 (\pm 0.0220)

Table 7: Results of CNN and CNN+LSTM based models; Source: Data compiled by author

Models	LogMeIn	TREC
CNN	0.916348 (\pm 0.0116)	0.813688 (\pm 0.0116)
CNN+LSTM	0.90359 (\pm 0.0119)	0.816875 (\pm 0.0094)

First, it can be observed that, as expected, the effects of the word embedding method are dependent on the dataset. Indeed, as results are not necessarily conclusive on the LogMeIn data, they are on the TREC one. Here I assume two effects must be taken into account. First, the larger the size of the vocabulary (i.e., total unique words in Table 3), the higher the model can leverage pre-trained vectors for similar ratios of words found/total words. Also, the higher the noise in the dataset, the higher will be the variance in performance of the downstream model.

Furthermore, looking in Table 6, it is also striking that the improvement in accuracy induced by the use of pre-trained vectors is dependent on the downstream model used to tackle the classification task. Indeed, compared to a random initialization, using the pre-trained FastText vectors can improve the accuracy by up to 10.2% using a CNN and 8.9% using a CNN+LSTM. Also, the impact is even greater if the subsequent model is not ideally trained. Indeed, in Table 5, from a random initialization of dimension 300 to the use of Dong & Huang the accuracy is potentially jumping from 65.62% to 79.72%, a 14.1% gain. The figures found are higher than the ones that in (Dhingra et al., 2017). As pre-trained vectors already carry information when fuelled to the subsequent classification model, they enhance the performance of the classifier. However, the nature of the gain, whether linear or not, has not, as far as I know, been investigated. It could be investigated by studying the relation between the number of epochs and the relative gains by using pre-trained vectors.

My hypothesis is that the marginal gain of using pre-trained vectors is diminishing as the number of epochs increases.

Second, simply doubling the dimension of the word embedding does not change the performance of the classification task with random initialization. However, doubling the dimension, allows reducing the variance of downstream results as observed in Figure 18.

Third, using pre-trained vectors yields indeed better results over random initialization which can be observed on the TREC results as well. As the matter of fact, except for GloVe pre-trained vectors, all embedding methods give better results than random initialization.

Fourth, using subword information from the pre-trained vectors of FastText does not improve either the performance. Further investigations using an architecture that uses character-level information such as in (Xiao and Cho, 2016; Zhang et al., 2015) should be performed to investigate whether these models can leverage these subword features better. In addition, it can be observed that methods leveraging the CBOW algorithm such as Word2Vec and FastText outperform GloVe. Looking at a 2D projections of the TREC vocabulary generated from FastText pre-trained vectors using t-distributed stochastic neighbour embedding (T-SNE³²)

³²A technique used to reduce the dimensionality of the vectors while keeping some features. The implementation has been done using the scikit-learn library (<http://scikit-learn.org/stable/index.html>) and matplotlib (<https://matplotlib.org/>)




Figure 17: 95% Confidence intervals of the results of the models on the LogMeIn dataset; Source: Author's own representation




Figure 18: 95% Confidence intervals of the results of the models on the TREC dataset; Source: Author's own representation

in Figure 19, we can observe that FastText vectors are better at capturing semantic information as words with similar meanings are clustered. However looking at the GloVe projections in Figure 20, we observe that fewer clusters appear and that the projection is similar to a random initialization with however increased variance. As a consequence, this variance also spills over the variance of performance of the models.

Finally, my proposal does not show a statistical difference with Dong & Huang's algorithm. As shown in Table 8, the number of out-of-vocabulary words is relatively low and the effectiveness of both methods is therefore hard to evaluate as a few information is added by the algorithm.

Further tests on different datasets with a greater number of OOV words should be performed. Indeed, words whose initialization is not random due to pre-training on the

dataset include "gotomeeting", "gotowebinar" or "seminario" which do not help much on determining whether the review is audio or not (low representative value). Likewise, on the TREC dataset these words include "spielberg", "mozambique", "gould". In TREC dataset, totally missing words include numbers such as "1991", "1967" or "327", or words such as 'occam', 'rockettes', 'quetzalcoatl', 'khrushchev', On the LogMeIn dataset missing words also includes numbers such as "995" or "65", misspelled words such as "presentationbefore" or 'probleme', and words in another language such as "perfekt", "einfache" or "reiniciar". A suggestion to improve the performance on the LogMeIn dataset is to use a combination of pre-trained vectors from different language as the dataset includes samples in another language than English.




Figure 19: T-SNE projection of FastText pre-trained vectors for the TREC vocabulary. Two clusters are shown; one presenting words about time (bottom) and the other one with modal verbs (top); Source: Author’s own representation




Figure 20: TT-SNE projection of word vectors for the TREC vocabulary. On the left vectors from GloVe. On the right random initialization of dimension 300; Source: Author’s own representation

Table 8: Descriptive statistics about the words found using different embedding methods; Source: Data compiled by author

	LogMeIn	TREC
Total unique words	2786	6987
Found in Word2Vec	2504	6036
Found in Glove	2575	6814
Found in FastText	2603	6501
<hr/>		
Proposal		
Found in both Word2Vec and generated vectors	393	863
Found only in Word2Vec	2111	5173
Found only in generated vectors	10	25
Not found	272	926

6. Conclusion

“As machines become more and more efficient and perfect, so it will become clear that imperfection is the greatness of man.” Ernst Fischer

As the literature in deep learning is flourishing, so is the range of models and their application. In this thesis, I have first described two common architectures used for text classification tasks namely convolutional neural networks (CNNs) and recurrent neural networks (RNNs). I have compared their performance and training procedure on a text classification task and found that CNNs are easier to train and yield better results. Nevertheless, according to the literature review, hybrid models combining both architectures can yield better results. Through my benchmark, I could not verify this statement as I could not reach optimal performance through my implementation but could highlight that sensitive factors for RNNs include the gradient update function and the number of epochs. I could also show that they are computably more expensive to train.

Also, I have discussed ways to convert textual data into inputs that the aforementioned models can leverage to improve their performance. I have highlighted that the use of pre-trained vectors can increase by up to 10.2% the performance of the subsequent model. Concretely, I have found that methods that generate word vectors based on a Continuous Bag of Word (CBOW) algorithm such as Word2Vec or FastText yield better results than count-based methods such as GloVe. Moreover, after observing the empirical results, I have stated that this gain is probably diminishing and therefore not linear as the subsequent models become fine-tuned. This could be subject to further research to confirm or not my hypothesis. I could also confirm that the gain was dependent not only on the subsequent models but also on the dataset used.

Finally, I proposed an algorithm for these models to deal with words that are unknown with unfortunately inconclusive results. Further evaluations are necessary with datasets that include a higher proportion of unknown words with a higher representative value. The benchmark used was designed to assess models on a classification task and not sufficient to evaluate my proposal.

While this thesis has been narrowed down to classification tasks for qualitative analysis, the use of neural networks is broad ranging from autonomous cars to automatic trading. The same way economists embraced the development of differential calculus to expand their models; entrepreneurs leveraged the spreading of the internet to create new business models, I expect managers and researchers to incorporate big data analytics into their day-to-day activities to understand better the world around us. However, as demonstrated during the hearing of Mark Zuckerberg, Facebook's CEO, in front of the U.S. Congress about the Cambridge Analytical scandal, even policymakers do not have a sound understanding of the current capabilities of modern techniques despite growing concerns about machines taking over human

jobs and big data techniques hijacking democracy. I, therefore, call for a democratisation of programming languages and a sensitisation of machine learning techniques as tools to solve problems, but also about the issues they raise. As a consequence, I hope this work demystified the functioning of neural networks and could be used as a gate by business students, entrepreneurs, managers, and teachers to enter the machine learning world.

References

- Agrawal, S. and Awekar, A. Deep Learning for Detecting Cyberbullying Across Multiple Social Media Platforms, 2018. URL <https://arxiv.org/pdf/1801.06482.pdf>. Retrieved from.
- Alcantara, G. Empirical analysis of non-linear activation functions for Deep Neural Networks in classification tasks, 2017. URL <https://arxiv.org/pdf/1710.11272.pdf>. Retrieved from.
- Apté, C., Damerou, F., and Weiss, S. M. Automated learning of decision rules for text categorization. *ACM Transactions on Information Systems (TOIS)*, 12(3):233–251, 1994.
- Baeza-Yates, R. and Ribeiro-Neto, B. Modern information retrieval. New York, 1999. URL <https://doi.org/10.1080/14735789709366603>.
- Baroni, M., Dinu, G., and Kruszewski, G. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 238–247, 2014.
- Becker, S., Le Cun, Y., et al. Improving the convergence of back-propagation learning with second order methods. In *Proceedings of the 1988 connectionist models summer school*, pages 29–37, 1988.
- Bengio, Y., Ducharme, R., Vincent, P., and Jauvin, C. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- Berger, M. J. Large Scale Multi-label Text Classification with Semantic Word Vectors. Technical Report, 2014. URL <https://cs224d.stanford.edu/reports/BergerMark.pdf>. Retrieved from.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent dirichlet allocation. *Journal of machine learning research*, 3(Jan):993–1022, 2003.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. Enriching Word Vectors with Subword Information, 2016. URL <https://doi.org/1511.09249v1>.
- Bottou, L. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998.
- Bottou, L. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.
- Cambria, E., Poria, S., Gelbukh, A., and Thelwall, M. Sentiment analysis is a big suitcase. *IEEE Intelligent Systems*, 32(6):74–80, 2017.
- Caudill, M. Neural Networks Primer, Part 1. AI Expert (Vol. 2). [CL Publications], 1986. URL <https://dl.acm.org/citation.cfm?id=38295>. Retrieved from.
- Cawley, G. C., Talbot, N. L., and Girolami, M. Sparse multinomial logistic regression via bayesian l1 regularisation. In *Advances in neural information processing systems*, pages 209–216, 2007.
- Chen, D., Bolton, J., and Manning, C. D. A thorough examination of the cnn/daily mail reading comprehension task. *arXiv preprint arXiv:1606.02858*, 2016.
- Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Cleverdon, C. Optimizing convenient online access to bibliographic databases. *Information services and Use*, 4:37–47, 1984.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Collobert, R. and Weston, J. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM, 2008.
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12(Aug):2493–2537, 2011.
- Conneau, A., Schwenk, H., Barrault, L., and Lecun, Y. Very deep convolutional networks for natural language processing. *arXiv preprint arXiv:1606.01781*, 2, 2016.
- Dahl, G. E., Sainath, T. N., and Hinton, G. E. Improving deep neural networks for lvsr using rectified linear units and dropout. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 8609–8613. IEEE, 2013.
- Debole, F. and Sebastiani, F. Supervised team weighting for Automated Text Categorization. Istituto Di Scienza E Tecnologie dell'Informazione, (MI), pp. 784-788, 2003. URL <http://nmis.isti.cnr.it/sebastiani/Publications/NEMIS04.pdf>. Retrieved from.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. Indexing by latent semantic analysis. *Journal of the American society for information science*, 41(6):391–407, 1990.
- Dhingra, B., Liu, H., Salakhutdinov, R., and Cohen, W. W. A comparative study of word embeddings for reading comprehension. *arXiv preprint arXiv:1703.00993*, 2017.
- Ding, Z., Xia, R., Yu, J., Li, X., and Yang, J. Densely connected bidirectional lstm with applications to sentence classification. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 278–287. Springer, 2018.
- Dong, J. and Huang, J. Enhance word representation for out-of-vocabulary on ubuntu dialogue corpus. *arXiv preprint arXiv:1802.02614*, 2018.
- Dozat, T. Incorporating Nesterov Momentum into Adam. ICLR Workshop (1), 2013–2016, 2016. Retrieved from.
- Economist. The world's most valuable resource is no longer oil, but data. May 6th, 2017. URL <https://www.economist.com/news/leaders/21721656-data-economy-demands-new-approach-antitrust-rules-worlds-most-valuable-resource>. Retrieved from.
- Elfwing, S., Uchibe, E., and Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.
- Elman, J. L. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.
- Erhan, D., Manzagol, P.-A., Bengio, Y., Bengio, S., and Vincent, P. The difficulty of training deep architectures and the effect of unsupervised pre-training. In *Artificial Intelligence and Statistics*, pages 153–160, 2009.
- Faraz, A. An elaboration of text categorization and automatic text classification through mathematical and graphical modelling. *Computer Science & Engineering: An International Journal (CSEIJ)*, 5(2/3):1–11, 2015.
- Forman, G. An extensive empirical study of feature selection metrics for text classification. *Journal of machine learning research*, 3(Mar):1289–1305, 2003.
- Frakes, W. B. Information retrieval, data structures & algorithms. Prentice Hall, 1992. URL <https://users.dcc.uchile.cl/~rbaeza/iradssbook/irbook.html>. Retrieved from.
- Fuhr, N., Hartmann, S., Lustig, G., Schwantner, M., Tzeras, K., and Knorz, G. Air/x: A rule-based multistage indexing system for large subject fields. In *Intelligent Text and Image Handling-Volume 2*, pages 606–623. LE CENTRE DE HAUTES ETUDES INTERNATIONALES D'INFORMATIQUE DOCUMENTAIRE, 1991.
- Genuer, R., Poggi, J.-M., and Tuleau-Malot, C. Variable selection using random forests. *Pattern Recognition Letters*, 31(14):2225–2236, 2010.
- Georgakopoulos, S. V., Tasoulis, S. K., Vrahatis, A. G., and Plagianakos, V. P. Convolutional neural networks for toxic comment classification. In *Proceedings of the 10th Hellenic Conference on Artificial Intelligence*, page 35. ACM, 2018.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- Goldberg, Y. A Primer on Neural Networks Models for Natural Language Processing 1–76, 2015. URL <https://doi.org/10.1613/jair.4992>.
- Gomez, J. C., Boiy, E., and Moens, M.-F. Highly discriminative statistical features for email classification. *Knowledge and information systems*, 31(1):23–53, 2012.
- Goodfellow, I., Bengio, Y., and Courville, A. Deep Learning. The MIT Press, 2016. URL http://www.deeplearningbook.org/front_matter.pdf. Retrieved from.
- Gövert, N., Lalmas, M., and Fuhr, N. A probabilistic description-oriented approach for categorizing web documents. In *Proceedings of the eighth international conference on Information and knowledge management*, pages 475–482. ACM, 1999.
- Greene, D. and Cunningham, P. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *Proceedings of the 23rd international conference on Machine learning*, pages 377–384. ACM, 2006.
- Guyon, I. and Elisseeff, A. An introduction to variable and feature selection.

- Journal of machine learning research*, 3(Mar):1157–1182, 2003.
- Hahnloser, R. H., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J., and Seung, H. S. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789):947, 2000.
- Harris, Z. S. Distributional structure. *Papers on Syntax*, pages 3–22, 1981.
- Hashimoto, K. and Tsuruoka, Y. Adaptive joint learning of compositional and non-compositional phrase embeddings. *arXiv preprint arXiv:1603.06067*, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Hinton, G., McClelland, J. L., and Rumerlhart, D. E. Distributed representations. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 1, 1986.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- IBM. What is big data?, 2018. URL <https://www-01.ibm.com/software/in/data/bigdata/>. Retrieved February 14, 2018.
- Janocha, K. and Czarnecki, W. M. On loss functions for deep neural networks in classification. *Classification. Schedae Informaticae*, 25:1–10, 2017.
- Jashki, M.-A., Makki, M., Bagheri, E., and Ghorbani, A. A. An iterative hybrid filter-wrapper approach to feature selection for document clustering. In *Canadian Conference on Artificial Intelligence*, pages 74–85. Springer, 2009.
- John, G. H., Kohavi, R., and Pfleger, K. Irrelevant features and the subset selection problem. In *Machine Learning Proceedings 1994*, pages 121–129. Elsevier, 1994.
- Johnson, R. and Zhang, T. Deep pyramid convolutional neural networks for text categorization. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 562–570, 2017.
- Kalchbrenner, N., Grefenstette, E., and Blunsom, P. A convolutional neural network for modelling sentences. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 655–665, 2014.
- Kaur, G. and Kaur, P. ISSN: 2454-132X Impact factor: 4.295 Review on Text Classification by NLP Approaches with Machine Learning and Data Mining Approaches. *International Journal of Advance Research , Ideas and Innovations in Technology*, 3, 767–771, 2017. Retrieved from www.ijariit.com.
- Khan, A., Baharudin, B., Lee, L. H., and Khan, K. A review of machine learning algorithms for text-documents classification. *Journal of advances in information technology*, 1(1):4–20, 2010.
- Kim, Y. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic gradient descent. *arXiv preprint arXiv:1412.6980*, 2015.
- Kira, K. and Rendell, L. A. A practical approach to feature selection. In *Machine Learning Proceedings 1992*, pages 249–256. Elsevier, 1992.
- Kiwiel, K. C. Convergence and efficiency of subgradient methods for quasi-convex minimization. *Mathematical programming*, 90(1):1–25, 2001.
- Kohavi, R. and John, G. H. Wrappers for feature subset selection. *Artificial intelligence*, 97(1-2):273–324, 1997.
- Kusner, M., Sun, Y., Kolkin, N., and Weinberger, K. From word embeddings to document distances. In *International Conference on Machine Learning*, pages 957–966, 2015.
- Le, H., Pham, Q., Sahoo, D., and Hoi, S. C. Urlnet: Learning a url representation with deep learning for malicious url detection. *arXiv preprint arXiv:1802.03162*, 2018.
- Lee, J. Y. and Dernoncourt, F. Sequential short-text classification with recurrent and convolutional neural networks. *arXiv preprint arXiv:1603.03827*, 2016.
- Lewis, D. D. An evaluation of phrasal and clustered representations on a text categorization task. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 37–50. ACM, 1992.
- Li, X. and Roth, D. Learning question classifiers: the role of semantic information. *Natural Language Engineering*, 12(3):229–249, 2006.
- Liu, P., Qiu, X., and Huang, X. Recurrent neural network for text classification with multi-task learning. *arXiv preprint arXiv:1605.05101*, 2016.
- Livnat, A., Papadimitriou, C., Pippenger, N., and Feldman, M. W. Sex, mixability, and modularity. *Proceedings of the National Academy of Sciences*, 107(4):1452–1457, 2010.
- Ma, Y., Peng, H., Khan, T., Cambria, E., and Hussain, A. Sentic lstm: a hybrid network for targeted aspect-based sentiment analysis. *Cognitive Computation*, 10(4):639–650, 2018.
- Mikolov, T., Kopecny, J., Burget, L., Glembek, O., et al. Neural network based language models for highly inflective languages. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4725–4728. IEEE, 2009.
- Mikolov, T., Karafiát, M., Burget, L., Černocký, J., and Khudanpur, S. Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*, pages 1045–1048, 2010.
- Mikolov, T., Kombrink, S., Burget, L., Černocký, J., and Khudanpur, S. Extensions of recurrent neural network language model. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5528–5531. IEEE, 2011.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., and Joulin, A. Advances in pre-training distributed word representations. *arXiv preprint arXiv:1712.09405*, 2017.
- Murty, M. R., Murthy, J., and PVGD, P. R. Text document classification based on least square support vector machines with singular value decomposition. *International Journal of Computer Applications*, 27(7):21–26, 2011.
- Ng, A. Y. Feature selection, l_1 vs. l_2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78. ACM, 2004.
- Nikam, S. S. A comparative study of classification techniques in data mining algorithms. *Oriental journal of computer science & technology*, 8(1):13–19, 2015.
- Olah, C. Understanding LSTM Networks – colah’s blog, 2017. URL <https://colah.github.io/posts/2015-08-Understanding-LSTMs/>. Retrieved April 14, 2018.
- Pennington, J., Socher, R., and Manning, C. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- Penrose, R. A generalized inverse for matrices. In *Mathematical proceedings of the Cambridge philosophical society*, volume 51, pages 406–413. Cambridge University Press, 1955.
- Porter, M. F. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
- Rahimian, V. and Ramsin, R. Designing an agile methodology for mobile software development: A hybrid method engineering approach. In *2008 Second International Conference on Research Challenges in Information Science*, pages 337–342. IEEE, 2008.
- Ramachandran, P., Zoph, B., and Le, Q. V. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- Ranganathan, V. and Natarajan, S. A new backpropagation algorithm without gradient descent. *arXiv preprint arXiv:1802.00027*, 2018.
- Robertson, S. E. and Walker, S. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR’94*, pages 232–241. Springer, 1994.
- Rong, X. word2vec parameter learning explained. *arXiv preprint arXiv:1411.2738*, 2014.
- Ruder, S. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016.
- Rumelhart, D. E., Hinton, G. E., Williams, R. J., et al. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1, 1988.
- Saeys, Y., Abeel, T., and Van de Peer, Y. Robust feature selection using ensemble feature selection techniques. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 313–325. Springer, 2008.
- Sag, I. A., Baldwin, T., Bond, F., Copestake, A., and Flickinger, D. Multiword expressions: A pain in the neck for nlp. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 1–15. Springer, 2002.
- Salinca, A. Convolutional neural networks for sentiment classification on

- business reviews. *arXiv preprint arXiv:1710.05978*, 2017.
- Salton, G., Wong, A., and Yang, C.-S. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.
- Sánchez, D., Martín-Bautista, M. J., Blanco, I., and de la Torre, C. J. Text knowledge mining: an alternative to text data mining. In *2008 IEEE International Conference on Data Mining Workshops*, pages 664–672. IEEE, 2008.
- Sebastiani, F. Machine learning in automated text categorization. *ACM computing surveys (CSUR)*, 34(1):1–47, 2002.
- Shen, Y., Huang, P.-S., Gao, J., and Chen, W. Reasonet: Learning to stop reading in machine comprehension. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1047–1055. ACM, 2017.
- Socher, R., Lin, C. C., Manning, C., and Ng, A. Y. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 129–136, 2011.
- Song, F., Liu, S., and Yang, J. A comparative study on text representation schemes in text categorization. *Pattern analysis and applications*, 8(1-2): 199–209, 2005.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Sundermeyer, M., Alkhouli, T., Wuebker, J., and Ney, H. Translation modeling with bidirectional recurrent neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 14–25, 2014.
- Sundström, J. Sentiment analysis of swedish reviews and transfer learning using convolutional neural networks, 2017. URL <http://uu.diva-portal.org/smash/get/diva2:1174477/FULLTEXT01.pdf>. Retrieved from.
- Sutskever, I., Martens, J., Dahl, G. E., and Hinton, G. E. On the importance of initialization and momentum in deep learning. *ICML (3)*, 28(1139-1147):5, 2013.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Tieleman, T. and Hinton, G. Lecture 6.5-rmsprop, coursera: Neural networks for machine learning. *University of Toronto, Technical Report*, 2012.
- Tsvetkov, Y., Faruqui, M., Ling, W., Lample, G., and Dyer, C. Evaluation of word vector representations by subspace alignment. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2049–2054, 2015.
- Turney, P. D. and Pantel, P. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37:141–188, 2010.
- Tzeras, K. and Hartmann, S. Automatic indexing based on bayesian inference networks. In *Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 22–35. ACM, 1993.
- Upadhyay, S., Chang, K.-W., Taddy, M., Kalai, A., and Zou, J. Beyond bilingual: Multi-sense word embeddings using multilingual context. *arXiv preprint arXiv:1706.08160*, 2017.
- Wang, B., Wang, L., Wei, Q., and Liu, L. Textzoo, a new benchmark for reconsidering text classification. *arXiv preprint arXiv:1802.03656*, 2018.
- Warde-Farley, D., Goodfellow, I. J., Courville, A., and Bengio, Y. An empirical analysis of dropout in piecewise linear networks. *arXiv preprint arXiv:1312.6197*, 2013.
- Weston, J., Bengio, S., and Usunier, N. Scaling up to large vocabulary image annotation. In *Proceedings of the 22st International Joint Conference on Artificial Intelligence*, pages 2764–2770, 2010.
- Wolfram, D. and Zhang, J. The influence of indexing practices and weighting algorithms on document spaces. *Journal of the American Society for Information Science and Technology*, 59(1):3–11, 2008.
- Wolpert, D. H. The lack of a priori distinctions between learning algorithms. *Neural computation*, 8(7):1341–1390, 1996.
- Xiao, Y. and Cho, K. Efficient character-level document classification by combining convolution and recurrent layers. *arXiv preprint arXiv:1602.00367*, 2016.
- Yin, W., Kann, K., Yu, M., and Schütze, H. Comparative study of cnn and rnn for natural language processing. *arXiv preprint arXiv:1702.01923*, 2017.
- Young, T., Hazarika, D., Poria, S., and Cambria, E. Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
- Yu, K., Liu, Y., Schwing, A. G., and Peng, J. Fast and accurate text classification: Skimming, rereading and early stopping. 2018.
- Yu, M. and Dredze, M. Learning composition models for phrase embeddings. *Transactions of the Association for Computational Linguistics*, 3:227–242, 2015.
- Zeiler, M. D. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- Zhang, X., Zhao, J., and LeCun, Y. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.
- Zhou, C., Sun, C., Liu, Z., and Lau, F. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*, 2015a.
- Zhou, C., Sun, C., Liu, Z., and Lau, F. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*, 2015b.
- Zhou, P., Qi, Z., Zheng, S., Xu, J., Bao, H., and Xu, B. Text classification improved by integrating bidirectional lstm with two-dimensional max pooling. *arXiv preprint arXiv:1611.06639*, 2016a.
- Zhou, Z., Zhu, X., He, Z., and Qu, Y. Question classification based on hybrid neural networks. In *2016 4th International Conference on Electrical & Electronics Engineering and Computer Science (ICEECS 2016)*. Atlantis Press, 2016b.



The Tax System and Corporate Payout Policies

Nicholas Herold

WHU - Otto Beisheim School of Management

Abstract

In this thesis, I examine how corporate taxes, dividend taxes, personal income taxes, and consumption taxes affect corporate payout behaviour. Using rich international panel data that consist of 40,609 firms across 115 countries from 1999 to 2013, I run linear regressions of each of the four tax rates on three payout variables which measure frequency and magnitude of regular cash dividends distributed by firms. In my baseline model, I find that the predictions of the new view – one of the two views in neoclassical theory – on short-run payout responses only partially hold true. Inconsistent with initial hypotheses, corporate taxes on average do not impact a firm's dividend payout behaviour in the short run. Regarding dividend taxes, my results show that the hypothesised dividend tax neutrality only holds true for the relative amount of dividends but not for a firm's likelihood to distribute, increase, and initiate dividends. Consistent with initial hypotheses, personal income taxes and consumption taxes trigger mostly large payout responses in terms of frequency and magnitude of dividend payouts. In my two model extensions, in which I focus on payout behaviour of cash-rich firms and employ a more flexible definition of the time horizon characterising short-run payout, my findings are again only partially in line with predictions of the new view on short-run payout responses. With these results, this thesis not only analyses well-investigated tax rates – corporate taxes and dividend taxes – for which current literature shows mixed empirical evidence but also examines hitherto scarcely considered tax rates – personal income taxes and consumption taxes – in the neoclassical framework and determines their impact on corporate payout.

Keywords: corporate payout; corporate tax; dividend tax; personal income tax; consumption tax

1. Introduction

Corporate payout policy is a fundamental part of corporate finance decisions besides deciding where to invest and how to finance projects of a firm. Taxes, however, reduce shareholders' wealth on both the firm level (e.g., via corporate taxes) and shareholder level (e.g., via dividend taxes), and thus likely distort payout decisions (Jacob and Jacob, 2013b). Hence, it is important for managers, shareholders,

and policy makers to understand how taxes affect corporate payout.

Previous literature has stipulated a variety of models showing whether and how a change in certain taxes potentially impacts payout decisions of firms. The most prominent frameworks in tax literature are neoclassical models which are typically divided into the old view (e.g., Harberger, 1962; Poterba and Summers, 1984) and the new view (e.g., Auerbach, 1979; King, 1977) suggesting that payout behaviour differs across firms due to different marginal sources of finance. Beyond neoclassical theories, agency models (e.g., Chetty and Saez, 2010; Jensen, 1986) provide an alternative explanation of how firms are predicted to react to changes in tax rates by considering the presence of agency issues. Neoclassical and agency models, however, mainly focus on corporate taxes and dividend taxes which both also constitute the primary area of interest in empirical studies as several tax reforms allowed a thorough examination of the

I would like to express my very great appreciation to Prof. Dr. Martin Jacob, my first supervisor, for his extensive guidance and structured advice during this research project. In particular, I would like to thank him for proposing the topic "The Tax System and Corporate Payout Policies", providing a pre-edited sample of core data, and engaging in fruitful discussions with me that altogether helped me successfully write this thesis. I would also like to express my deep gratitude to Dr. Antonio De Vito, my second supervisor, for his far-reaching support through insightful deep-dive sessions on specialised statistical software and meaningful intellectual discussions that shaped the outcome of this research project. Finally, I would like to thank Frédéric Herold, Ralf Herold, and Jan Spörer for their helpful comments and suggestions that improved the quality of this thesis.

impact of corporate taxes¹ and, in particular, dividend taxes on corporate payout. In the setting of dividend tax cuts in the U.S. in 2003 (e.g., Chetty and Saez, 2005) and in Sweden in 2006 (Jacob and Michaely, 2017), a variety of studies support predictions of agency models², but empirical evidence on the neoclassical predictions remains heavily disputed³. Tax research also discusses the impact of personal income taxes on corporate payout, but the exact definition varies strongly⁴ and rarely refers to taxes on labour income in the context of dividends or share repurchases⁵. Regarding consumption taxes, previous literature has hitherto solely examined the effect on corporate investment (Jacob et al., 2018) without considering the effect on corporate payout. This is surprising given that, intuitively, corporate payout is somehow related to the level of investment since managers can either (i) immediately invest earnings in projects and distribute resulting profits in future periods or (ii) immediately distribute earnings to shareholders or (iii) retain earnings for future investments and payout.

Due to the different state of literature across tax rates, this thesis aims at providing a comprehensive overview of how a change in corporate taxes, dividend taxes, personal income taxes, and consumption taxes affects payout decisions of firms. Specifically, this thesis contributes to contemporary literature in two ways. First, it adds to the ongoing discussion about mixed empirical evidence on neoclassical theories for well-researched tax rates (i.e., corporate taxes, dividend taxes). Second, it bridges the current gap in literature by embedding scarcely considered tax rates (i.e., personal income taxes, consumption taxes) in the neoclassical frameworks and investigating their impact on corporate payout. To achieve this, I use international panel data

¹In the context of corporate taxes, Poterba et al. (1987), for instance, examines how the 1986 Tax Reform Act in the U.S. is predicted to lower corporate savings and reduce tax incentives to retain earnings and distribute dividends.

²These studies conclude that frictions such as agency issues (Chetty and Saez, 2005, Jacob and Michaely, 2017) and shareholder conflicts (Jacob and Michaely, 2017) reduce the responsiveness of corporate payout in case of a dividend tax change.

³Chetty and Saez (2005) argue that listed U.S. firms responded to the 2003 dividend tax cut in accordance with the old view. By contrast, Brav et al. (2008) conclude that the immediate payout response of these firms was only temporary and that the dividend tax cut was of “second-order importance ... [as only] firms ‘sitting on the fence’ [to initiate dividends]” (p.390) were primarily affected.

⁴Wu (1996), for example, uses the term “personal taxes” (p.293) synonymously for dividend taxes in his empirical study on the payout behaviour of listed U.S. firms. Likewise, Lewellen and Lewellen (2006) employ “personal tax rates on interest, dividends, and realized capital gains” (p.5) in their single- and multi-period models when theorising how corporate payout changes depending on the firm’s source of finance.

⁵In the context of private firms, Jacob and Michaely (2017), for instance, argue that the taxation of labour income vis-à-vis dividends incentivises only a specific group of owners to adjust the corporate payout of their firm due to “strong empirical evidence that, with a limited number of owner[-managers in closely-held corporations], there is strong substitutability between dividends and wages (the other possible form of payout to owners in private firms)” (p.3219). Other empirical studies also examine the sole impact of personal income taxes whose scope, however, is mainly on macroeconomic variables such as economic growth (Gale and Samwick, 2016; Palić et al., 2017).

with focus on non-financial, non-utility, non-transportation, and non-telecommunication firms across 115 countries over the period 1999 to 2013 with sufficient variation in tax rate changes. My estimation strategy involves three steps: (i) Pre-analysis, (ii) baseline regression, and (iii) extensions to the baseline model. Inspired by Jacob et al. (2018), the pre-analysis is mainly based on a linear probability model to rule out the concern that tax rate changes are determined by macroeconomic factors. The baseline regression is the main analysis in this thesis where I investigate the average effect on corporate payout in the same year in which a change in one of the four tax rates occurs. Consistent with previous studies (e.g., Jacob and Jacob, 2013a), I measure payout, which is defined as regular cash dividends due to insufficient data on other payout channels, by three dependent variables covering frequency and relative amounts of dividends. Beyond the baseline model, I also introduce two extensions which consider heterogeneity in payout responses potentially caused by different levels of cash holdings (e.g., Jacob and Michaely, 2017) and the impact of tax rate changes on payout one year after a tax rate change occurs.

The results of my baseline regression show that the average payout response only partially follows neoclassical predictions on short-run payout responses as stipulated by the new view. Inconsistent with initial expectations, corporate taxes on average do not change a firm’s dividend payout behaviour in the year where a tax change becomes effective. Similarly, the hypothesised “dividend tax neutrality” (Chetty and Saez, 2010, p.5) only holds with respect to the relative amount of dividends. Vice versa, a change in dividend taxes interestingly impacts a firm’s propensity to pay dividends and likelihood to increase or initiate dividends in different directions (i.e., sign of coefficients differs) even though the relative effect size is small. Personal income taxes show mostly significant coefficients suggesting that a higher tax rate increases the attractiveness of investments in corporate projects such that firms invest more. Thus, they exhibit a slightly lower propensity to pay dividends and distribute considerably lower amounts in the short run. The results on consumption taxes are fully in line with my initial hypotheses implying that a rise in this tax rate increases the tax wedge (Jacob et al., 2018) exerting pressure on profits of corporate projects such that firms invest less in the short run and therefore distribute, increase, and initiate dividends more frequently and pay higher relative amounts.

The baseline extensions reveal mostly similar findings. Cash-rich firms appear to react more strongly compared to the average payout response in terms of their likelihood to increase or initiate dividends if personal income taxes, consumption taxes, and (depending on the fixed effect) corporate taxes are changed. Although the payout response of cash-rich firms is expected to match more closely short-run predictions of the new view, the results do not fully confirm this expectation and thus are again only partially in line with predictions of neoclassical theory. When considering the payout response one year after a tax rate change, corporate taxes again do not appear to impact payout behaviour on average.

Also, corporate payout is mostly not neutral to a change in dividend taxes. Interestingly, the coefficient of personal income taxes on a firm's likelihood to increase or initiate dividends changes its sign suggesting that payout decisions in subsequent periods are increasingly determined by the fact that firms bear higher labour costs from an increase in this tax rate. Regarding consumption taxes, the results are very similar to the findings of the baseline model.

The remaining part of this thesis is divided into seven further sections. Section 2 provides a profound theoretical background on both neoclassical frameworks old view and new view which I use as a foundation to formulate hypotheses on how each of the four tax rates affects dividend payout. Section 3 presents my methodology and displays descriptive statistics on all variables of interest employed in the main analysis. In section 4, I conduct my pre-analysis using the linear probability model and test whether my dataset contains sufficient variation in tax rate changes. Sections 5 and 6 show the results of my baseline regression and extensions to the baseline model, respectively. In section 7, I test for robustness of my baseline results. Finally, the conclusion of this thesis is shown in section 8.

2. Theoretical Background and Hypothesis Formulation

Even though various theories provide explanations on how taxes might affect corporate payout decisions, empirical studies mostly analyse their findings in the two neoclassical frameworks: The old view (Feldstein, 1970; Harberger, 1962, Harberger, 1966; Poterba, 2004; Poterba and Summers, 1984) and the new view (Auerbach, 1979; Auerbach and Hassett, 2003; Bradford, 1981; King, 1977). Conceptually, these views differ in the underlying assumption of how firms fund the additional project (i.e., what constitutes a firm's marginal source of finance). That is, the old view assumes that firms finance new projects via new equity whereas the new view is built on the idea that retained earnings are used (see also Chetty and Saez, 2005). In the following, old view and new view will be incorporated into an intuitive single-period model based on previous studies (Alstadsæter et al., 2017; Becker et al., 2013; Chetty and Saez, 2010; Lewellen and Lewellen, 2006) to illustrate the effect of corporate taxes (τ^C), dividend taxes (τ^{Div}), personal income taxes (τ^I), and consumption taxes (τ^{VAT}) on corporate payout decisions. Figure 1 visualises how an increase in each of these tax variables impacts investors' after-tax returns which, in turn, changes investment and payout decisions. For simplicity, my hypotheses are built on two assumptions. First, I restrict corporate payout to regular cash dividends and abstract from special dividends and share buybacks⁶. Second, I

⁶As discussed by Chetty and Saez (2005), firms have three payout channels: Regular cash dividends, special dividends, and share buybacks. I exclude special dividends as they occur infrequently and are difficult to measure such that clear causal inference would not be possible. I also exclude share buybacks since my dataset does not contain any information on this payout channel. However, I acknowledge the increasing importance of share

use a highly stylised definition of τ^I in my hypothesis formulation which involves both "personal taxes on interest" (Lewellen and Lewellen, 2006, p.5) and personal taxes on labour income⁷.

2.1. Old View and New View in the Single-Period Model

In the old view, the individual investor decides at the beginning of period t whether to (i) invest in the firm's project by buying new equity or (ii) invest in an alternative investment opportunity which is for simplicity assumed to be a risk-free bond (see also Alstadsæter et al., 2017). If the investor decides to invest \$1 in a firm's project (see arrow A in Figure 1), the project will generate profits depicted by the pre-tax rate of return, r . These profits are assumed to be distributed in form of dividends in $t+1$ and are subject to double taxation due to taxes levied on both the firm level and the shareholder level (Jacob and Jacob, 2013b). On the firm level, corporate taxes are levied on pre-tax project earnings. Assuming that firms fully distribute their after-tax profits as dividends at the beginning of period $t+1$, potential payout $\$1[1+r]$ is effectively reduced to actual payout (i.e., gross dividends distributed by firms) $\$1[1+r(1-\tau^C)]$ (arrow B). On the shareholder level, these dividends are further reduced by dividend taxes finally yielding the after-tax dividend income (i.e., net dividends received by shareholders) $\$1[1+r(1-\tau^C)(1-\tau^{Div})]$ (arrow C). By contrast, the alternative investment in a risk-free bond generates an interest payment denoted by the coupon rate, i , and is not subject to double taxation. In this scenario, only personal income taxes reduce pre-tax interest income $\$1[1+i]$ to the level of after-tax interest income $\$1[1+i(1-\tau^I)]$ (arrow D). Thus, the rational investor will always invest in the firm's project if and only if the after-tax dividend income (arrow C) is larger than after-tax income on the bond (arrow D). Hence, the investor invests in the firm if pre-tax return on the firm's project, r , at least meets the individual investor's minimum required rate of return, r_{old}^* , which is defined as the pre-tax rate of return on the firm's project where the investor is indifferent between buying new equity and investing in the risk-free bond in t .

In the new view, the firm decides at the beginning of period t whether to (i) invest in a profit-generating project and

buybacks as an alternative payout channel (Chetty and Saez, 2005; Von Eijie and Megginson, 2008). In this context, Jacob and Jacob (2013b) have shown that the relative taxation of dividends vis-à-vis capital gains matters for a firm's payout channel choice. If capital gains are taxed at a higher rate than dividends, firms would prefer distributing dividends over share buybacks and vice versa, as this yields a higher after-tax income for shareholders. Thus, share buybacks and the corresponding relative taxation should be incorporated in future studies.

⁷This treatment is in line with the current tax code of the United States (Office of the Law Revision Council, 2018). However, it does not hold for other tax jurisdictions such as Germany where interest income (25% flat tax) is presently taxed at a different rate compared to labour income (45% top marginal income tax rate) (Federal Ministry of Justice and Consumer Protection, 2017). In the empirical part of this thesis, I nonetheless try to interpret my results using the hypothesised mechanism of my simplified single-period model, but I also acknowledge that the definition of τ^I varies across countries and therefore add footnote 28 on this topic in section 5.

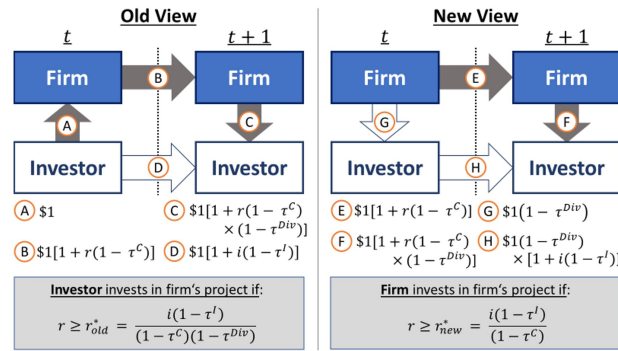


Figure 1: Investment and Payout Decisions in the Old View and the New View

subsequently distribute dividends at the beginning of t+1 or (ii) directly distribute its retained earnings at the beginning of period t to shareholders who invest in a risk-free bond immediately after receiving this dividend payment in t. Similar to the old view, the project will generate a pre-tax return, r , if the firm decides to invest \$1 in the project, and r will be again diminished by corporate taxes and dividend taxes (arrows E and F) yielding the shareholder's after-tax dividend income $\$1[1+r(1-\tau^c)(1-\tau^{Div})]$ at the beginning of period t+1. If the firm decides not to invest in its project, the dividends distributed in period t are again subject to τ^{Div} yielding net dividends $\$1(1-\tau^{Div})$ (arrow G). After investing these net dividends in a risk-free bond, investors finally obtain $\$1(1-\tau^{Div}) \times [1+i(1-\tau^I)]$ (arrow H). Assuming that firms aim at maximising shareholders' after-tax wealth, the firm will invest in its project if and only if the shareholder's after-tax dividend income in t+1 (arrow F) is larger than the after-tax income on the risk-free bond (arrow H). Likewise, r must again at least meet the individual investor's minimum required rate of return, r_{new}^* , such that the firm invests in its project instead of directly distributing dividends in period t.

2.2. Hypothesis Formulation

Based on this theoretical foundation, four hypotheses will be outlined in the following. These hypotheses aim at explaining the potential effect of each of the four taxes on dividend payout in the light of both the old view and the new view, and thus consider that the marginal source of finance impacts dividend payout at different points in time: Firms in the old view can only distribute dividends in period t+1 (i.e., they receive new equity in period t which they invest in new projects generating profits and thus dividends of the next period) whereas firms in the new view can decide whether to distribute dividends in period t (i.e., immediate payout) or pay dividends in period t+1 (i.e., from profits generated by project investment in period t). A summary on the hypothesised effects of an increase in taxes on dividends is shown in Table 1.

Hypothesis 1: In t+1, an increase in corporate taxes decreases dividends in both old view and the new view. In t, an increase in corporate taxes increases dividends in the new view.

If τ^c increases, firms for which new equity is the marginal source of finance are expected to pay lower dividends in period t+1 (Chetty and Saez, 2010). Ceteris paribus, higher corporate taxes increase the individual investor's minimum required rate of return, r_{old}^* , as investors demand a higher pre-tax return on projects, r , to receive the same after-tax dividend income as if taxes did not change. In other words, investing in firms becomes less attractive relative to investing in a risk-free bond since the after-tax returns on the bond (arrow D) remain unaffected; corporate after-tax earnings (arrow B) and the shareholder's after-tax dividend income (arrow C), however, decrease. Thus, fewer projects can offer an r that meets r_{old}^* of investors such that more investors decide not to buy new equity in period t. As investors invest in fewer projects, firms generate lower profits, and therefore dividend payout is expected to decrease in t+1.

If firms predominantly finance their projects via retained earnings, an increase in τ^c is predicted to increase dividends in period t but decrease dividends in period t+1 (Chetty and Saez, 2010).

Similar to the old view, higher corporate taxes in the new view increase r_{new}^* while r itself remains unaffected. Thus, firms are expected to distribute dividends in period t to maximise after-tax wealth of investors instead of investing in profitable projects for which r is below the higher r_{new}^* . Firms will therefore invest in fewer projects leading to lower profits for firms. This, in turn, results in lower dividends to be distributed in t+1. As firms, however, decide whether to invest in corporate projects or directly pay out dividends to shareholders in period t, a lower level of investments in t directly corresponds to higher dividends in t.

Hypothesis 2: In t+1, an increase in dividend taxes is expected to decrease dividends in the old view while the new view predicts no change in dividends in t and t+1.

In the old view, an increase in τ^{Div} is expected to result in lower dividends in t+1 (e.g., Jacob and Jacob, 2013b). The line of argumentation is similar to the effect of τ^c on dividends predicted by the old view: A rise in τ^{Div} increases r_{old}^* , fewer projects with their given r will be able to satisfy the higher r_{old}^* , investors invest less in corporate projects, fewer

Table 1: Effect of an Increase in Tax Rates on Corporate Payout

This table shows the effect of an increase in corporate taxes (column (1)), dividend taxes (column (2)), personal income taxes (column (3)), and consumption taxes (column (4)) on a firm’s dividend payout in periods t and t+1 as predicted by the old view and the new view.

		Increase in Tax Rate			
		τ^C	τ^{Div}	τ^I	τ^{VAT}
		(1)	(2)	(3)	(4)
Old View	t+1	↓ Div	↓ Div	Direct: ↑ Div Indirect: ↓ Div	↓ Div
	t	↑ Div	No Change	Direct: ↓ Div Indirect: ↑ Div	↑ Div
New View	t+1	↓ Div	No Change	Direct: ↑ Div Indirect: ↓ Div	↓ Div

projects are realised, and firms generate lower profits resulting in a lower level of dividends in t+1.

The new view stipulates “dividend tax neutrality” (Chetty and Saez, 2010, p.5) which implies that a rise in τ^{Div} has no effect on a firm’s dividend payout decision. If τ^{Div} is increased at the beginning of period t and remains at this new level until the end of period t+1, net dividends received by the investor in t (arrow G) or t+1 (arrow F) would be equally reduced. Consequently, r_{new}^* stays constant and the firm’s decision to distribute dividends in t or invest in a project followed by paying dividends in t+1 is not impacted at all. In essence, the new view expects dividend payout in t and t+1 to remain unaffected if τ^{Div} changes. This prediction is likely to hold in the absence of agency issues and shareholder conflicts⁸.

Hypothesis 3: An increase in personal income taxes reveals an ambiguous effect on corporate payout in both old view (t+1) and new view (t, t+1).

Irrespective of the marginal source of finance, an increase in τ^I impacts dividend payouts in two ways. First, there is a direct effect on the after-tax returns on the bond (old view: arrow D; new view: arrow H). An increase in τ^I reduces these after-tax returns such that investing in corporate projects becomes relatively more attractive for the investor (old view) and the firm (new view). In other words, an increase in τ^I reduces the investor’s minimum required rate of return, r_{old}^* and r_{new}^* . Consequently, investors (old view) and firms (new view) will invest more in corporate projects in t leading to

more projects being realised, and higher profits generated by firms which, in turn, result in higher dividends in t+1 in both old view and new view. The new view additionally predicts an effect on dividends in period t. More investments in corporate projects in t automatically mean that less retained earnings are available to be distributed in t. Hence, dividends in t are expected to decline if τ^I increases.

Second, there is an indirect effect on the project pre-tax returns, r , which are a function of τ^I . Intuitively, a rise in τ^I increases labour costs of firms. Assuming that revenues generated by projects remain constant, this rise in labour costs decreases r leading to lower after-tax earnings on the firm level (old view: arrow B; new view: arrow E) and reduced net dividends in t+1 (old view: arrow C; new view: arrow F). Hence, fewer projects will be able to meet r_{old}^* and r_{new}^* such that investors (old view) and firms (new view) invest less in corporate projects in t resulting in lower profits and a lower level of dividends in t+1 in both neoclassical models. Once again, the new view additionally predicts an effect on dividends in t. A lower level of investments in corporate projects in t directly corresponds to more retained earnings which can be distributed in period t. Thus, dividend payments in t are expected to rise if τ^I increases. This hypothesis is likely to hold if workers have a strong negotiation power vis-à-vis firms, for example in the presence of strong unions, allowing workers to shift part of the tax burden to firms (Alesina et al., 2002).

Hypothesis 4: In t+1, an increase in consumption taxes decreases dividends in both old view and new view. In t, an increase in consumption taxes increases dividends in the new view.

Similar to personal income taxes, consumption taxes have an indirect effect on corporate payout. “Consumption taxes drive a wedge between the price that consumers pay and the price that producers receive. Hence, [the] firms’ profitability is expected to decrease when consumption taxes increase” (Jacob et al., 2018, p.3). In other words, an increase in τ^{VAT} lowers the pre-tax return on firms’ projects, r . Thus, fewer projects are able to meet r_{old}^* and r_{new}^* such that investors (old

⁸For simplicity, I abstract from agency issues. However, I acknowledge that governance plays an important role in corporate payout decisions. In the setting of the 2003 dividend tax cut in the U.S., Chetty and Saez (2005) show that agency issues shape payout responses as well-governed firms (i.e., firms with strong principals such as institutional investors with large shareholdings) or agents whose interests are aligned with shareholders’ interests (e.g., due to high executive share ownership) respond more strongly to a tax cut in dividends. Likewise, Jacob and Michaely (2017) find that agency issues and shareholder conflicts mute a firm’s payout response in the context of the 2006 dividend tax cut in Sweden.

view) and firms (new view) invest less in corporate projects in t . In $t+1$, this yields lower profits and therefore lower dividends according to both old view and new view. In addition, the new view stipulates higher dividends in t as a lower level of investment in corporate projects means that more retained earnings will be distributed in t .

3. Data and Descriptive Statistics

The majority of data used in my analysis was issued by the WHU Chair of Business Taxation which, in turn, withdrew these data from three main sources. First, firm-level information on listed firms around the world over the period 1997 to 2013 was derived from the Compustat North America and Global database. Second, annual tax rates involving corporate taxes, dividend taxes, personal income taxes, and consumption taxes were retrieved from tax handbooks released by Ernst & Young, KPMG, PricewaterhouseCoopers, and Deloitte. Third, country-level statistics comprising macroeconomic variables, country governance indicators, income group descriptions⁹, and region group classifications¹⁰ were extracted from the World Bank database.

After consolidating all data¹¹, I converted each monetary variable which was originally quoted in each firm's local currency into USD using average annual exchange rates provided by the WHU Chair of Business Taxation. Subsequently, I conducted general data cleaning by excluding firms with SIC codes 4000-4999 and 6000-6999¹². The general data

⁹The original dataset provided by the WHU Chair of Business Taxation contained some missing data entries on income group descriptions which, however, were required to successfully change fixed effects in the robustness section. Using World Bank data, I manually amended 15 income group descriptions in total for Argentina, Jamaica, New Zealand, and Nigeria where some country-years contained a missing entry. Please refer to the Excel file *WorldBank_Data_Income_History* stored on the USB device for details on the missing income group descriptions for these four countries. Furthermore, I retrieved the full historical income group dataset from the World Bank database covering the period 1998 to 2013 for 66 countries like Estonia, Saudi-Arabia, Taiwan, and Vietnam for which firm data already existed but no information on income groups was present. Please refer to the Excel file *WorldBank_income_group_history_missing* stored on the USB device for details on the missing income group descriptions for these 66 countries.

¹⁰I extracted region names and region codes from the World Bank database and added these data to the information provided by the WHU Chair of Business Taxation. This step was required to cluster all countries in my dataset by region and successfully make changes to the definition of my fixed effects in the robustness section. Please refer to the Excel file *WorldBank_Data_Region_Codes* stored on the USB device for detailed region information provided by the World Bank.

¹¹The WHU Chair of Business Taxation additionally provided data on Tobin's q with high coverage across firms in my sample which I merged into my dataset. The initially provided dataset revealed a poor coverage of Market Value (i.e., market value of equity) and thus Tobin's q . Other attempts to generate Market Value via Common Shares Outstanding and Price Close (i.e., market price per share) hardly increased the coverage.

¹²This treatment is similar to Chetty and Saez (2005) and Jacob and Jacob (2013b) dropping financial firms (6000-6999) and utility firms (4900-4999) because firms in these industries are subject to "additional regulations and hence might have different payout behaviour" (Chetty and Saez, 2005, p.798). I additionally excluded transportation and (tele-)communication firms (4000-4899) since most of these firms are privatised companies which are small in number but contribute disproportionately much to aggregate dividends especially in the European Union (Von Eije and Megginson, 2008).

cleaning was further complemented by dropping all observations which appeared illogical for my analysis in six steps. First, I dropped observations for which there was no information on total assets or when total assets were negative. Second, I removed bankrupt firms (i.e., firms with a book value of common equity equal to or lower than zero) from my dataset. Third, I dropped firms with negative values for cash and short-term investments, sales, and cash dividends¹³. Fourth, firms with leverage values smaller than zero and larger than or equal to one were excluded, too. Fifth, I also removed observations with negative tax rates or tax rates exceeding one. Sixth, I excluded negative values for the macroeconomic variables GDP per Capita, Openness, Government Debt, and Interest Payments which, realistically, are not smaller than zero.

Lastly, I winsorised my lagged firm controls and non-dummy dependent variables below the 1st percentile and above the 99th percentile of observations to mitigate biased results caused by large outliers. After all adjustments, the sample used for my baseline regression consists of 42,672 firms across 115 countries over the period 1997 to 2013¹⁴. Table 2 shows descriptive statistics for all dependent variables, tax rates, firm-level variables, and country-level variables contained in this sample.

4. Pre-Analysis: Variation in Tax Rate Changes and Linear Probability Model

Prior to running a baseline regression, two major concerns have to be addressed. First, the underlying sample has to overcome the frequently objected "lack of compelling tax variations" (Chetty and Saez, 2005, p. 792) to avoid a small number of events potentially biasing my results. Otherwise, it would be difficult to make a well-founded generalisation of the impact of taxes on corporate payout. Second, all four tax rates, which constitute the independent variables of my baseline regression, have to be exogenous to conduct convincing causal inference.

To address the first concern, the sample of my baseline regression indeed contains sufficient variation in all four tax rates. Across all 115 countries over the period 1999 to 2013, there are 315 corporate tax changes (48 increases; 267 decreases), 144 dividend tax changes (72 increases; 72 decreases), 217 personal income tax changes (76 increases; 141 decreases), and 105 consumption tax changes (72 increases; 33 decreases).

To address the second concern, I employ a linear probability model inspired by Jacob et al. (2018) to determine likely country-level correlates with the magnitude of tax rate changes and ideally rule out issues "that tax policy is not

¹³Cash dividends refer to the variable Cash Dividends (Cash Flow) in my consolidated dataset serving as a proxy for corporate payout.

¹⁴Since all tax rates start in 1999, my baseline sample effectively starts in 1999, too. Consequently, the number of firms used in my baseline regression drops to 40,609 while the number of countries remains unchanged.

Table 2: Summary Statistics of Main Variables

This table is an overview of summary statistics of my main variables covering 42,672 firms across 115 countries over the period 1997 to 2013. Panel A shows the three payout variables which are used as dependent variables in my baseline regression. Panel B presents the four tax variables of interest. Panel C and Panel D depict firm-level and country-level controls, respectively. Please see table A.1 in the appendix for detailed definitions of all main variables. Note: Summary statistics of Dividend Yield (t) in Panel A and all firm-level controls in Panel C are based on the winsorised version of the respective variables to debias the mean.

Variable	N	Mean	Standard Deviation	25th percentile	Median	75th percentile
Panel A: Payout Variables						
Dividend Payer (t)	272,182	0.6584	0.4742	0	1	1
Dividend Increase (t)	224,464	0.2458	0.4305	0	0	0
Dividend Yield (t)	251,472	0.0087	0.0230	0.0000	0.0002	0.0052
Panel B: Tax Variables						
Corporate Tax	345,995	0.3215	0.0742	0.2700	0.3300	0.3900
Dividend Tax	345,374	0.1965	0.1183	0.1000	0.2000	0.2643
Personal Income Tax	345,374	0.3973	0.0938	0.3500	0.4000	0.4641
Consumption Tax	325,902	0.1073	0.0627	0.0519	0.1000	0.1700
Panel C: Firm-level Controls						
Leverage	369,79	0.0933	0.1564	0.0007	0.0112	0.1167
Cash Holdings (L. TA)	338,23	0.1270	0.2728	0.0020	0.0203	0.1130
Cash Flow	327,475	0.0110	0.1602	-0.0002	0.0029	0.0541
Profits	337,816	0.0267	0.2107	-0.0037	0.0517	0.1141
Retained Earnings	336,703	-0.2718	1.3584	-0.0032	0.0033	0.0581
Ln(Sales Growth)	323,767	0.0876	0.4336	-0.0460	0.0730	0.2182
Tobin's q	279,478	1.5000	3.5913	0.3319	0.6838	1.4020
Firm Size	388,244	6.5550	3.0444	4.3861	6.3837	8.4699
Panel D: Country-level Controls						
Macroeconomic Variables						
Ln(GDP per Capita)	363,858	9.6126	1.3841	8.6600	10.4301	10.5557
GDP Growth	363,943	3.5811	3.4686	1.7292	3.1400	5.1472
Inflation	363,943	2.7069	4.3171	0.8477	2.0327	3.7157
Openness	304,225	0.7265	0.8648	0.2829	0.4831	0.6549
Deficit	269,554	-2.6677	3.9788	-4.8523	-3.1779	0.0177
Interest Payments	279,996	0.0225	0.0123	0.0150	0.0230	0.0276
Government Debt	196,656	60.9354	37.7064	40.0881	53.5029	64.0318
Governance Indicators						
Voice and Accountability	371,063	0.6718	0.8952	0.3900	1.0100	1.3500
Political Stability	371,058	0.3317	0.8166	-0.2000	0.6000	0.9600
Government Effectiveness	371,047	1.1320	0.7665	0.4000	1.4600	1.7500
Regulatory Quality	371,047	0.9723	0.7836	0.4200	1.1900	1.6200
Rule of Law	371,063	0.9889	0.7968	0.2900	1.3300	1.6100
Control of Corruption	371,047	0.9796	0.9781	0.0500	1.2900	1.8350

exogenously determined [sic] but related to changes in economic conditions" (Jacob et al., 2018, p.15). Results of the linear probability model are presented in Table 3.

Overall, changes in dividend taxes and personal income taxes appear to be exogenous. Changes in corporate taxes and consumption taxes, however, are likely to be influenced by the macroeconomic factors GDP Growth and Ln(GDP per

Capita) and the factors GDP Growth and Deficit, respectively. The significance of GDP Growth suggests that, based on my dataset, policy makers tend to decrease (increase) corporate taxes and consumption taxes in periods where the economy is in a boom phase (recession). Also, corporate taxes are likely to be increased (decreased) if a country generates a higher (lower) level of GDP per Capita implying that an economy

Table 3: Results of Linear Probability Model

This table presents how macroeconomic determinants of tax rates potentially affect the magnitude of a tax rate change in corporate taxes (column (1)), dividend taxes (column (2)), personal income taxes (column (3)), and consumption taxes (column (4)). The definitions of all tax rates and macroeconomic variables are outlined in the appendix in Table A.1. I include country fixed effects and region-year fixed effects in all four regressions. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Magnitude of Tax Rate Change in			
	Corporate Taxes (1)	Dividend Taxes (2)	Personal Income Taxes (3)	Consumption Taxes (4)
GDP Growth	-0.0006** (0.0002)	-0.0004 (0.0008)	0.0000 (0.0006)	-0.0006*** (0.0001)
Ln(GDP per Capita)	0.0258** (0.0113)	-0.0353 (0.0292)	0.0114 (0.0218)	-0.0003 (0.0052)
Inflation	-0.0000 (0.0001)	0.0002 (0.0003)	0.0000 (0.0002)	-0.0000 (0.0001)
Deficit	-0.0002 (0.0002)	0.0005 (0.0008)	-0.0005 (0.0004)	0.0004* (0.0002)
Openness	0.0089 (0.0072)	0.0105 (0.0215)	-0.0024 (0.0187)	0.0037 (0.0036)
Interest Payments	0.1823 (0.1348)	0.0757 (0.3864)	0.1096 (0.1674)	0.1045 (0.0632)
Observations	800	743	743	709
Country FE	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes
Adjusted R-squared	0.014	-0.072	-0.106	0.095

becomes more (less) productive and thus wealthier (poorer). Consumption taxes are likely to rise (be reduced) if a country's budget deficit increases (decreases). This result seems to be reasonable as, intuitively, an increase in government spending needs to be somehow financed; this finding, however, is not in line with the linear probability model results of [Jacob et al. \(2018\)](#) despite using (almost) the same underlying dataset¹⁵. Across all tax rates, the variables Inflation, Openness, and Interest Payments appear to be insignificant.

Based on these results, I include a GDP-Growth-Ln(GDP per Capita) cluster in a fixed effect used in my baseline regression¹⁶. This way, it is possible to account for potential endogeneity in tax rate changes and compare countries which are economically similar in terms of GDP level and GDP growth rates.

¹⁵Interestingly, my linear probability model results on consumption taxes (column (4)) are based on 709 observations whereas [Jacob et al. \(2018\)](#) rely on 664 observations.

¹⁶I deliberately excluded Deficit from the fixed effect because the linear probability model only proves marginal significance of this variable ($p=.091$). This stands in stark contrast to GDP Growth ($p=.024$; $p=.000$) and Ln(GDP per Capita) ($p=.027$). Thus, the significance of Deficit arguably could have emerged by chance. Also, excluding Deficit is unlikely to adversely affect my baseline results since it is correlated with the other two macroeconomic variables incorporated in the fixed effect. Please refer to the correlation matrix in 2.0_LPM_RESULTS_(EDITED) for further details.

5. Baseline Regression

To investigate the average effect of a tax rate change on corporate payout, I stipulate the following linear regression model using the ordinary least squares (OLS) estimation method:

$$\begin{aligned} \text{Payout}_{i,j,t} = & \alpha_0 + \beta_1 \text{CorporateTax}_{j,t} + \beta_2 \text{DividendTax}_{j,t} \\ & + \beta_3 \text{PersonalIncomeTax}_{j,t} + \beta_4 \text{ConsumptionTax}_{j,t} \\ & + \delta_1 \Phi_{i,j,t-1} + \delta_2 \Gamma_{j,t} + \alpha_i + \alpha_{g,k,t} + \epsilon_{i,j,t} \end{aligned} \quad (1)$$

The dependent variable $\text{Payout}_{i,j,t}$ is a payout measure of firm i headquartered in country j in year t . This payout measure is a placeholder for the three payout variables Dividend Payer (t), Dividend Increase (t)¹⁷ and Dividend Yield (t)¹⁸

¹⁷Dividend Increase (t) covers a firm's likelihood to substantially increase (if a firm was a dividend payer in year $t-1$) or initiate dividends in year t (if a firm was no dividend payer in year $t-1$). This variable is particularly interesting as "against the background of the general stickiness of dividends ..., the decision to initiate or substantially increase dividends is a strong commitment to a long stream of cash outlays (as opposed to a simple 1-year commitment that can be easily reversed)" ([Jacob and Jacob, 2013b, p.1256](#)). In my baseline regression, a substantial increase in dividends is defined as an increase by at least 25%. This might be viewed as sufficiently strict since the number of observations where firms pay dividends in year $t-1$ and increase them in year t drops from 90,546 to 48,002 while observations covering initiations remain unaffected.

¹⁸Dividend Yield (t) is defined as the dividend-to-total-assets ratio similar

based on previous literature (e.g., [Jacob and Jacob, 2013a](#); [Alstadsæter et al., 2017](#)). All variable definitions are presented in the appendix in Table A.1. In my baseline model, I restrict all dependent variables to the time identifier (t) representing year t (i.e., the year in which a change in taxes first becomes effective) for two reasons. First, most firms are likely to react quite fast to a change in taxes to maximise profits and shareholder value. This assumption seems to be reasonable as most tax rate changes are announced several months or, in favourable cases, a year in advance prior to becoming effective. Second, [Brav et al. \(2008\)](#) have shown that tax-related payout motives gain importance in the immediate aftermath of a tax rate change but only play a minor role in subsequent periods.

The independent variables of interest are the four tax variables $\text{CorporateTax}_{j,t}$, $\text{DividendTax}_{j,t}$, $\text{PersonalIncomeTax}_{j,t}$, and $\text{ConsumptionTax}_{j,t}$. The baseline regression also includes two control vectors to account for alternative determinants of corporate payout on the firm level and the country level which are denoted by $\Phi_{i,j,t-1}$ and $\Gamma_{j,t}$, respectively. Control vector $\Phi_{i,j,t-1}$ consists of the following eight firm-level controls which are frequently used in literature¹⁹ (e.g., [Jacob and Jacob, 2013b](#)): Leverage, Cash Holdings (L. TA), Cash Flow, Profits, Retained Earnings, Tobin's q, Sales Growth, and Firm Size²⁰. To rule out endogeneity concerns, I additionally lag each variable included in this vector by one year. Control vector $\Gamma_{j,t}$ consists of nine country-level controls. Inspired by [Jacob et al. \(2018\)](#), I included the three macroeconomic variables GDP Growth, Ln(GDP per Capita), and Inflation²¹ besides the six governance indicators Voice and Accountability, Political Stability, Government Effectiveness,

to [Alstadsæter et al. \(2017\)](#). Due to poor coverage of market capitalisation, conventional definitions such as "the dollar amount of dividends paid out in year t+1 divided by the end-of-year t equity market value" ([Jacob and Jacob, 2013a](#), p.1251) are not used.

¹⁹My baseline regression does not include any proxy for ownership structure although previous literature has shown that it heavily impacts corporate payout decisions as agency issues (e.g., [Chetty and Saez, 2005](#)) and shareholder conflicts might arise or be mitigated ([Jacob and Michaely, 2017](#)). Due to lack of compelling data, however, I cannot proxy for ownership structure (e.g., via percentage of closely held shares ([Jacob and Jacob, 2013b](#))) which may reduce the explanatory power of my model.

²⁰Leverage considers that creditors in firms with a high debt-to-capital ratio tend to urge these firms to refrain from distributing dividends (e.g., [Jensen, 1986](#)). Cash Holdings (L. TA) acknowledges that cash-rich firms, intuitively, have more funds to be distributed to shareholders (e.g., [Chetty and Saez, 2010](#)). I incorporate Cash Flow to capture the positive effect of a company's cash flow on dividends (e.g., [Jacob and Jacob, 2013a](#)) which goes beyond considering pure cash holdings. Profits are a proxy for internal resources in addition to cash holdings (e.g., [Jacob and Michaely, 2017](#)). Retained Earnings acknowledge that mature firms tend to have larger retained earnings which are more likely to distribute dividends and pay larger amounts (e.g., [Jacob and Jacob, 2013b](#)). Tobin's q is "a proxy for stock undervaluation and growth opportunities" ([Jacob and Jacob, 2013b](#), p.1254) and Sales Growth also measures growth opportunities (e.g., [Alstadsæter et al., 2017](#)). Firm Size is used since larger firms, intuitively, have a higher propensity to pay dividends and distribute larger amounts (e.g., [Jacob and Michaely, 2017](#)).

²¹I exclude the remaining macroeconomic variables Openness, Deficit, Interest Payments, and Government Debt from my baseline regression due to poor coverage which could potentially bias my results. Please refer to Table 1 showing that these four variables have a considerably lower coverage than

Regulatory Quality, Rule of Law, and Control of Corruption in this control vector.

I employ two fixed effects in my baseline regression. First, I use firm fixed effects, α_i , to control for firm characteristics which potentially impact payout decisions (e.g., firm age). Second, I employ group-industry-year fixed effects, $\alpha_{g,k,t}$, where group (subscript g) refers to a GDP-Growth-Ln(GDP per Capita) cluster. This cluster is additionally combined with a specific industry k in year t to compare firms in the same industry-year which also operate in economically similar countries in terms of GDP level and GDP growth rates. Finally, I use heteroskedasticity-robust standard errors which are clustered at the country level since firms headquartered in country j are exposed to the same tax system.

Returning to the four hypotheses in section 2, my baseline regression only allows a clear causal interpretation of results with respect to the new view in period t because of two reasons. First, periods t and t+1 in the single-period model are a simplified theoretical abstraction where period t models short-run effects (i.e., payout responses in year t and year t+1) and period t+1 models long-run effects (i.e., payout responses in more distant future periods such as year t+5). Long-run effects, in particular, are difficult to measure since period t+1 might represent many years (e.g., ten years) until old-view firms eventually start distributing dividends. Similarly, it could take new-view firms a long time until they show a payout response matching predictions of period t+1 assuming that no other tax rate change occurs in the meantime. Also, dividend payout in more distant future periods is increasingly determined by confounding factors (e.g., a firm's financial performance and general economic developments). Therefore, I measure payout in year t (baseline regression) and year t+1 (second baseline extension) and thus restrict the interpretation of my results to short-run responses matching period t in the neoclassical models. Second, the sole consideration of short-run responses, by definition, only allows validation or rejection of new-view predictions in period t as the old view does not predict any payout response in period t (i.e., old-view firms receive new equity and thus cannot adjust their payout behaviour to a change in taxes in period t). This argumentation is further supported when considering a typical old-view firm characterised by young age, high growth rates, and financial constraints ([Chetty and Saez, 2010](#)) suggesting that they are less likely to distribute dividends in order to grow further. If these firms become more mature, grow at lower rates, and have sufficiently high financial reserves, they are more likely to distribute dividends on a regular basis²² (see also [Sinn, 1991](#)). Thus, even though my sample likely consists of dividend-paying firms which might

GDP Growth, Ln(GDP per Capita), and Inflation. Nonetheless, I incorporate Openness, Deficit, and Interest Payments into the vector $\Gamma_{j,t}$ as a robustness test in section 7.

²²Consistent with the model of [Sinn \(1991\)](#), I assume that old-view firms transform into new-view firms over time. This assumption is supported by [DeAngelo et al. \(2006, p. 227\)](#): "Consistent with a life-cycle theory of dividends, the fraction of publicly traded . . . firms that pay dividends is high when retained earnings are a large portion of total equity (and of total as-

exhibit some characteristics of old-view firms (e.g., financial constraints), it is still reasonable to focus on the new view when interpreting my baseline results.

Based on my initial hypotheses in the light of the new view in period t , I derive the following four predictions. First, I expect the coefficient of $\text{CorporateTax}_{j,t}$ to be significant and positive (i.e., $\beta_1 > 0$) across all payout variables as an increase in corporate taxes in year t exerts pressure on firms which are financed via retained earnings to directly distribute dividends in year t instead of investing in a project whose after-tax returns, and thus dividends in future periods, decline from higher corporate taxes²³. Second, I predict the variable $\text{DividendTax}_{j,t}$ to be insignificant as implied by the hypothesised “dividend tax neutrality” (Chetty and Saez, 2010, p.5). Third, the effect of a change in personal income taxes on corporate payout depends on whether the direct effect or the indirect effect prevails. The direct effect predicts lower (higher) dividends in year t in terms of probability and magnitude due to higher (lower) attractiveness of investing in corporate projects compared to other investment opportunities (e.g., bonds). Conversely, the indirect effect forecasts higher (lower) dividends in year t due to higher (lower) labour costs yielding a lower (higher) the relative attractiveness of corporate projects. Thus, if the direct (indirect) effect dominates, I expect $\text{PersonalIncomeTax}_{j,t}$ to have significant and negative (positive) coefficients across all payout variables (i.e., direct: $\beta_3 < 0$; indirect: $\beta_3 > 0$). Fourth, I predict the coefficients of $\text{ConsumptionTax}_{j,t}$ to be significant and positive (i.e., $\beta_4 > 0$) across all dependent variables as a rise in consumption taxes increases the tax wedge (Jacob et al., 2018) which reduces corporate investment. This, in turn, makes more retained earnings available to be distributed as dividends in year t instead.

The compact version of my baseline results is shown in Table 4. Columns (1), (2), and (3) (columns (4), (5), and (6)) report the coefficients of each tax rate (the relative effect of a one-percentage-point increase in a tax rate) with regard to the dependent variables Dividend Payer (t), Dividend Increase (t), and Dividend Yield (t), respectively. A detailed results overview of coefficients (standard errors) for all regressors (i.e., including firm-level and country-level controls) is shown in Table A.2 in the appendix.

Interestingly, the results of my baseline regression only partially confirm my hypotheses based on the new view in period t . On average, corporate taxes do not seem to impact any payout variable due to insignificant coefficients in all three columns. This suggests that firms do statistically not

sets) and falls to near zero when most equity is contributed rather than earned.”

²³This explanation assumes one of the two following conditions. First, shareholders must be sufficiently strong to exert pressure on management teams. As shown by Chetty and Saez (2005), this is the case for firms with strong principals, i.e., large institutional investors such as pension funds and independent directors are major shareholders. Second, firms in which their management teams hold a high percentage of shares are more likely to act on behalf of their shareholders as managers are major shareholders themselves and thus benefit from higher dividends, too (Chetty and Saez, 2005).

respond to a change in corporate taxes which is not consistent with the new view in period t . Although the insignificant coefficients suggest that an effect is statistically not present, it is surprising that the sign of all coefficients is negative and not, as expected, positive. My hypothesis on corporate taxes does not predict this outcome which I therefore recommend examining in future studies.

According to my baseline regression, the hypothesised “dividend tax neutrality” (Chetty and Saez, 2010, p.5) only holds with regard to a firm’s relative amount of dividends²⁴ due to an insignificant coefficient in column (3). With respect to a firm’s propensity to pay and the likelihood to increase or initiate dividends, dividend taxes seem to influence a firm’s payout behaviour due to significant coefficients in columns (1) and (2). Surprisingly, the direction of the effect (i.e., sign of coefficient) differs between the dependent variables. Inconsistent with initial expectations, column (1) shows that a rise in dividend taxes in year t by one percentage point (in the following abbreviated as pp) increases the probability of a firm distributing dividends in year t by 0.24pp²⁵. The relative effect, however, is comparatively small as a one-pp increase in dividend taxes in year t increases the probability of a firm paying dividends in year t by 0.36%²⁶ relative to the average probability of a firm paying dividends. Despite this small relative effect size, neither the new view nor empirical studies evidencing a negative relation between dividend taxes and a firm’s propensity to pay dividends (e.g., Chetty and Saez, 2005) support the positive coefficient in column (1). Thus, the reason for this effect should be further investigated in future studies. On the contrary, a rise in dividend taxes in year t by one pp results in a lower likelihood to increase or initiate dividends in year t by 0.25pp. This result is again not in line with the new view but would be supported by the empirical

²⁴This interpretation appears to depend on the observations in my sample as the coefficient of $\text{DividendTax}_{j,t}$ in column (3) is only marginally not significant ($p=.109$). Therefore, it is possible that a slightly different sample composition could have shown significant results implying that the “dividend tax neutrality” (Chetty and Saez, 2010, p.5) does not hold. In such a scenario, the effect on Dividend Yield (t) would have a similar interpretation as the effect on Dividend Payer (t) (see column (1)), but the relative effect size of 1.02% is considerably larger given that Dividend Yield (t) is defined as dividends divided by lagged total assets. Please refer to footnote 26 for the relative effect calculation.

²⁵Please note that all coefficients in Table 4 pertain to tax rates ranging from 0 (i.e., 0%) to 1 (i.e., 100%) in my original dataset. For example, the tax rate 0.30 for a specific country-year refers to a tax rate equal to 30%. To interpret the coefficients as a one-pp increase in a tax rate (i.e., a tax rate change by one unit equivalent to one pp), I mathematically transform these tax rates into whole numbers ranging from 0 to 100 (i.e., I multiply these tax rates by 100) and simultaneously divide the respective coefficients by 100. Hence, the coefficient 0.2389 (0.24 after rounding) turns into 0.002389 (0.0024). As my dependent variables are also defined between 0 (0%) and 1 (100%), I can interpret a one-pp increase in the corporate tax rate (e.g., from 30% to 31%) as a change in my dependent variable by 0.2389pp (0.24pp).

²⁶The relative effect is calculated by dividing 0.002389 (i.e., transformed coefficient) by 0.6584 (i.e., the average value of the dependent variable; please refer to Table 2 presenting summary statistics on all dependent variables). This yields 0.0036 or 0.36%. Please refer to the tab Relative Effect Calculation in the Excel file 3.0_Baseline_Results_(Edited)_Final for all calculations.

Table 4: Results of Baseline Regression (incl. Relative Effects)

This table shows the compact version of the results of my baseline regression from 1999 to 2013. Additionally, the relative effect of a change in taxes on each dependent variable is included in columns (4), (5), and (6). Relative effects are computed by dividing the coefficient of a tax variable by the mean of the respective dependent variable. The dependent variables are Dividend Payer (t) (column (1) and (4)), Dividend Increase (t) (column (2) and (5)), and Dividend Yield (t) (column (3) and (6)). All independent variables are defined in the appendix in Table A.1. I include firm fixed effects and gdp-cluster-industry-year fixed effects in all three regressions. Please note that gdp-cluster is a placeholder representing a GDP-Growth-Ln(GDP per Capita) cluster. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Please refer to Table A.2 in the appendix for a more detailed overview (including number of observations, adjusted R-Squared, etc.) of my regression results. Note: The mean of Dividend Yield (t) is based on the winsorized version of the respective variable to avoid biased results due to the presence of extreme values.

	Coefficients			Relative Effects		
	(1)	(2)	(3)	(4)	(5)	(6)
Corporate Tax	-0.1549 (0.1821)	-0.2884 (0.1872)	-0.0154 (0.0166)	-0.24%	-1.17%	-1.78%
Dividend Tax	0.2389*** (0.0655)	-0.2461** (0.1213)	0.0088 (0.0054)	0.36%	-1.00%	1.02%
Personal Income Tax	-0.3154*** (0.1185)	-0.1049 (0.0903)	-0.0543*** (0.0129)	-0.48%	-0.43%	-6.27%
Consumption Tax	1.2559* (0.6998)	1.8259*** (0.6860)	0.0883** (0.0420)	1.91%	7.43%	10.19%

findings of [Chetty and Saez \(2005\)](#) showing that listed firms increasingly initiated or increased dividends in the six quarters following the dividend tax cut in the U.S. in 2003²⁷. Yet, the relative effect size is again small: A one-pp increase in dividend taxes merely reduces a firm's likelihood of increasing or initiating dividends in year t by 1.00% compared to the average likelihood of increasing or initiating dividends.

Regarding personal income taxes, the direct effect on payout appears to dominate the indirect effect with regard to a firm's propensity to pay dividends and the relative amount of dividends due to negative and significant coefficients in columns (1) and (3), respectively. Thus, firms appear to acknowledge that an increase in personal income taxes makes future dividends from corporate investments more attractive for investors who would otherwise invest in less attractive investments such as bonds after receiving dividends in year t. Therefore, firms invest more in year t, which, in turn, yields lower dividends in year t²⁸. To be more precise, a one-pp in-

crease in personal income taxes reduces a firm's propensity to pay dividends by 0.32pp and the amount of dividends distributed by 0.05pp of lagged total assets with small (i.e., negative 0.48%) and large (i.e., negative 6.27%) relative effect sizes, respectively. Due to a negative but insignificant coefficient in column (2), an increase in personal income taxes in year t, however, reveals that neither the direct effect nor the indirect effect eventually dominates in terms of a firm's likelihood to increase or initiate dividends in year t.

Due to significant coefficients in columns (1), (2), and (3), consumption taxes seem to impact all payout variables. Also, the direction of the effect is consistent with my hypothesis as all coefficients are positive. For example, a one-pp rise in consumption taxes yields a 1.26pp (1.83pp; 0.09pp) increase in a firm's propensity to pay dividends (likelihood to increase or initiate dividends; amount of dividends relative to lagged total assets). Also, columns (5) and (6) suggest that the relative effect size of a change in consumption taxes is moderately large (7.43%) and considerably large (10.19%) compared to the average likelihood to increase or initiate dividends and the average relative amount of dividends, respectively. By contrast, column (4) suggests that the relative effect size of a change in consumption taxes compared to the average likelihood to pay dividends at all is moderately small (1.91%). These results provide evidence that a rise in consumption taxes increases the tax wedge which, in turn, reduces corporate investment. Thus, more retained earnings are available to be distributed as dividends which results in a higher probability to pay, increase or initiate, and a larger relative amount of dividends distributed in year t.

To conclude, my hypotheses based on the new view in period t are only partially confirmed²⁹. My baseline results mostly corroborate the neoclassical predictions on personal

²⁷[Chetty and Saez \(2005\)](#), however, implicitly argue that the payout response measured over these six quarters in 2003 and 2004 is sufficient to validate long-run responses. Thus, they conclude that their results resemble predictions of the old view. As outlined above, however, long-run payout responses are technically difficult to measure in year t and year t+1. Hence, I would be cautious when considering the conclusion of [Chetty and Saez \(2005\)](#) and rather interpret my findings in the context of the new view in period t.

²⁸As mentioned in footnote 7, this interpretation assumes that firms and investors are in a tax jurisdiction where labour income and interest income are taxed at the same rate (e.g., the U.S.). Thus, my interpretation of the results at first glance seems to be vague when considering other countries. However, regarding the disproportionately high percentage of dividend-paying firms in my sample which are headquartered in the United States (17,786 out of 166,084 and 159,721 observations), they might have vastly contributed to this result due to major personal income tax changes in the U.S. in 2003 and 2013. Yet, there might also be an alternative explanation especially for other tax jurisdictions than the U.S. which I recommend examining in future studies.

²⁹I also test for three alternative thresholds defining a substantial in-

income taxes (i.e., direct effect mostly prevails) and consumption taxes (i.e., positive and significant coefficients across all payout variables). The neoclassical predictions on corporate taxes and dividend taxes, however, are mostly not supported by the baseline results. In the following, I therefore additionally test whether my hypotheses hold in the context of (a) cash-rich firms and (b) dependent variables with the new time identifier ($t+1$).

6. Extensions to the Baseline Model

6.1. Heterogeneity in Payout Responses due to Different Levels of Cash Holdings

The first extension of my baseline model considers heterogeneity in payout responses arising from different levels of cash holdings. As the average payout response only partially confirms my initial hypotheses, I disentangle the average response and consider the payout behaviour of cash-rich firms. According to neoclassical theory, cash-rich firms are predicted to follow the new view because these firms have sufficient cash holdings and retained earnings to finance new projects or distribute dividends (Chetty and Saez, 2010). To account for differences in cash holdings, I therefore define the dummy variable High Cash which is equal to one if a firm has a cash-to-total-assets ratio (Cash Holdings (TA)) larger than the median value³⁰ of this ratio in a given country-year. Subsequently, I interact each tax variable with High Cash to examine whether cash-rich firms exhibit a different payout response compared to the average response of my baseline regression.

I expect one of the two following outcomes to materialise. First, the response of cash-rich firms could match predictions of the new view in period t more closely than suggested by the average response in my baseline regression. In this case, I would expect the sign of the combined effect (i.e., average effect plus marginal effect if firm is cash rich) of each tax rate in this extension to have the same sign as the beta of the respective tax rate as originally predicted for the baseline regression. For instance, the combined effect of corporate taxes on dividends is expected to be positive if a firm is cash rich, i.e., the marginal effect is predicted to be positive and significant offsetting the negative average effect. Second, cash-rich

firms might simply react more strongly compared to the average response. That is, the sign of the interaction term coefficient is expected to be positive (negative) if the coefficient of the average response shows a positive (negative) sign. The latter expectation is based on the findings of Alstadsæter et al. (2017) showing that cash-rich firms respond more strongly to a tax cut in dividends. Building on this result, I extend the scope of Alstadsæter et al. (2017) and include three further tax rates.

Table 5 reports the results of my first baseline extension. In column (2), the coefficients of the interaction terms³¹ suggest that cash-rich firms react more strongly to a change in personal income taxes and consumption taxes. In fact, cash-rich firms exhibit an even higher and even lower likelihood of increasing or initiating dividends if personal income taxes and consumption taxes rise, respectively. To be more precise, cash-rich firms are 0.15pp less (0.16pp more) likely to increase or initiate dividends if personal income taxes (consumption taxes) increase by one pp which corresponds to a total decrease (increase) in a firm's likelihood to increase or initiate dividends by 0.17pp (1.92 pp) if the firm is cash rich. In relative terms, a change in personal income taxes and consumption taxes implies that cash-rich firms respond more strongly almost by factor 8 and by 9.08%³², respectively. However, a firm's propensity to pay dividends and the relative amount of dividends do not vary with different levels of cash holdings among firms.

From these results, I infer that the predictions of Alstadsæter et al. (2017) (i.e., my second expected outcome) conceptually hold for personal income taxes and consumption taxes with respect to a firm's likelihood to increase or initiate dividends. The expected stronger payout response of cash-rich firms to a change in dividend taxes, however, cannot be inferred from my results. Also, cash-rich firms do not appear to respond more strongly to changes in corporate taxes. Furthermore, the stronger response of cash-rich firms in the event of a change in personal income taxes (here: direct effect) and consumption taxes confirms my first expected outcome, too. In other words, the combined effect (i.e., average effect plus marginal effect if firm is cash rich) of both tax rates in this extension has the same sign as the beta of the respective tax rate as originally predicted for the baseline regression. This suggests that cash-rich firms as a proxy

crease for the variable Dividend Increase (t): 10%, 50%, and 100%. When modifying this threshold, results are similar for $\text{DividendTax}_{j,t}$ and $\text{PersonalIncomeTax}_{j,t}$ in significance and magnitude. Results on other tax rates, however, vary depending on the specification. Please refer to the Excel file 3.1_DivIncr(t)_THRESHOLDS_(EDITED) for detailed regression results on all alternative threshold definitions.

³⁰I define the median value by country-year instead of country-industry-year to rule out a potentially incorrect High Cash classification of firms. For example, firms operating in cash-rich industries would be classified as cash-poor firms if they have lower cash holdings compared to their industry peers. This would occur even though these below-median firms have significantly larger cash holdings compared to firms in cash-poor industries. Hence, I abstract from industry-specific differences in cash reserves and solely acknowledge that cash holdings might vary across countries.

³¹Please note that I only interpret the interaction effects since I am interested in whether cash-rich firms exhibit a different payout response compared to the average response. Consequently, I disregard the average effects in this regression as they are not examined by my research question in this section.

³²These numbers describe by how much more strongly cash-rich firms react relative to the average effect. Therefore, I use the coefficients in column (2) of Table 5 and divide the combined effect (i.e., average effect plus marginal effect if firm is cash rich) by the average effect and finally subtract 1. A change in personal income taxes yields a stronger response of cash-rich firms by factor 7.9319 or 793.19% (i.e., $[((-0.1515) + (-0.0191)) / (-0.0191) - 1]$ which is the same as the marginal effect dividend by the average effect, i.e., $[(-0.1515) / (-0.0191)]$). Similarly, a change in consumption taxes causes a stronger response by 9.08% (i.e., $0.0908 = 0.1594 / 1.7563$) if the firm is cash rich.

Table 5: Differences in Payout Behaviour due to Different Cash Holdings

This table displays the regression results showing whether different levels of cash holdings explain different payout responses between firms from 1999 to 2013. I define Dividend Payer (t) (column (1) and (4)), Dividend Increase (t) (column (2) and (5)), and Dividend Yield (t) (column (3) and (6)) as my dependent variables. All independent variables are defined in the appendix in Table A.1. Additionally, I interact each tax rate with a dummy (High Cash) equal to one if a firm has a cash-over-total-assets ratio (Cash Holdings (TA)) larger than the median in a given country-year. In columns (1), (2), and (3), I include firm fixed effects and gdp-cluster-industry-year fixed effects. Please note that gdp-cluster is a placeholder representing a GDP-Growth-Ln(GDP per Capita) cluster. In columns (4), (5), and (6), I include firm fixed effects and country-industry-year fixed effects. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Corporate Tax	-0.1096 (0.1743)	-0.2414 (0.1888)	-0.0111 (0.0165)			
Corporate Tax × High Cash	-0.0880 (0.0807)	-0.0553 (0.0535)	-0.0078 (0.0049)	-0.0731 (0.0857)	-0.0912** (0.0431)	-0.0066 (0.0047)
Dividend Tax	0.2632*** (0.0672)	-0.2383* (0.1241)	0.0111** (0.0053)			
Dividend Tax × High Cash	-0.0472 (0.0528)	-0.0127 (0.0141)	-0.0044 (0.0044)	-0.0483 (0.0529)	-0.0169 (0.0145)	-0.0045 (0.0044)
Personal Income Tax	-0.2852** (0.1181)	-0.0191 (0.0953)	-0.0539*** (0.0135)			
Personal Income Tax × High Cash	-0.0463 (0.0599)	-0.1515*** (0.0409)	-0.0001 (0.0041)	-0.0496 (0.0618)	-0.1247*** (0.0353)	-0.0009 (0.0039)
Consumption Tax	1.2762* (0.6956)	1.7563** (0.6766)	0.0852** (0.0415)			
Consumption Tax × High Cash	-0.0685 (0.0905)	0.1594** (0.0638)	0.0073 (0.0059)	-0.0705 (0.0980)	0.1424** (0.0589)	0.0072 (0.0058)
Observations	178,161	168,309	178,161	177,275	167,454	177,275
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
GDP-Cluster-Industry-Year FE	Yes	Yes	Yes	No	No	No
Country-Industry-Year FE	No	No	No	Yes	Yes	Yes
Adjusted R-squared	0.789	0.161	0.683	0.794	0.171	0.693

for new-view firms respond even more clearly in accordance with the new view in period t compared to the average response. Contrarily, the insignificant interaction coefficients of Corporate Tax and Dividend Tax with High Cash suggest that cash-rich firms respond statistically as strong as other firms to a change in these tax rates and thus do not corroborate predictions of the new view. This finding is surprising given that cash-rich firms in particular are predicted to follow the new view. As neoclassical theory does not explain this result, other factors might have contributed to this outcome or the interaction with High Cash does not proxy new-view firms sufficiently well.

The results are very similar if I choose different fixed effects to rule out the concern that unobservable characteristics in a certain country, specific industry, and a given year explain my results. I therefore replace the GDP-Growth-Ln(GDP per Capita)-cluster-industry-year fixed effect by country-industry-year fixed effects in columns (4), (5), and (6) “to absorb any previously omitted unobservable time-varying characteristics at the [country-industry] level” (Jacob et al., 2018, p.21). Similar to columns (1), (2), and (3), the level of significance and the magnitude of the interaction coefficients remain mostly unchanged. The sole difference is that

the interaction of Corporate Tax and High Cash becomes significant, too. Hence, I cannot fully rule out the concern that “unobservable country-(industry)-year variables [are] correlated with ... [the] tax changes” (Jacob et al., 2018, pp.21-22)³³. Yet, the negative coefficient of Corporate Tax and High Cash in column (5) is again not in line with predictions of the new view in period t which could therefore be an alley of future research.

6.2. Impact on Payout in Year t+1

The second extension of my baseline model considers the effect of tax rate changes in year t on payout in year t+1

³³Also, my results are not robust to different definitions of High Cash. I define cash-rich firms in two alternative ways: Firms have a level of cash holdings (Cash Holdings (TA)) such that they are in (a) the top tercile and (b) the top quartile of cash holdings in a given country-year. Across both alternative definitions, interaction terms which show significant coefficients when using the median as the High Cash threshold are not significant anymore (and vice versa). Also, the dependent variables which are impacted by a tax rate change differ depending on the High Cash threshold. Changing the fixed effects also reveals an unclear picture of whether cash-rich firms react significantly differently than the average response. Thus, the effect of taxes on the response of cash-rich firms remains unclear.

Table 6: Results on Payout in Year t+1

This table presents the regression results showing how a change in taxes affects corporate payout in the year after a tax rate change (i.e., year t+1) from 1999 to 2013. I define Dividend Payer (t+1) (column (1)), Dividend Increase (t+1) (column (2)), and Dividend Yield (t+1) (column (3)) as my dependent variables. All independent variables are defined in the appendix in Table A.1. I include firm fixed effects and gdp-cluster-industry-year fixed effects in all three regressions. Please note that gdp-cluster is a placeholder representing a GDP-Growth-Ln(GDP per Capita) cluster. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
Corporate Tax	-0.1552 (0.1464)	0.0925 (0.2794)	-0.0084 (0.0104)
Dividend Tax	0.2296*** (0.0612)	0.0173 (0.1275)	0.0164*** (0.0041)
Personal Income Tax	-0.1693 (0.1252)	0.5854*** (0.1371)	-0.0330*** (0.0070)
Consumption Tax	1.1282* (0.5684)	-0.0726 (0.8565)	0.0861** (0.0366)
Observations	142,493	137,004	142,493
Controls	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
GDP-Cluster-Industry-Year FE	Yes	Yes	Yes
Adjusted R-squared	0.805	0.150	0.709

assuming that tax rates in a certain country are not altered every year (i.e., tax regimes are quite stable). Generally, investigating the impact on payout in year t+1 seems to be reasonable as some firms might respond to a tax rate change with a certain delay (i.e., not in year t already) which, for example, might depend on how much in advance a tax rate change is announced before becoming effective and how flexibly individual firms are able to react. As outlined in section 5, dependent variables with time identifier (t+1) also cover short-run payout responses and are therefore expected to be in line with predictions of the new view in period t.

Table 6 reports the results of this second baseline extension. Interestingly, the results are quite similar to the ones of the baseline model suggesting that a tax rate change in year t does not only impact payout in year t but also has an effect on payout in year t+1. For instance, a change in the corporate tax rate on average again does not affect corporate payout. Surprisingly, the coefficient in column (2) becomes positive and the coefficient in column (3) converges to zero. This matches predictions of the new view in period t more closely than in the baseline regression but cannot be fully corroborated due to insignificant coefficients. Furthermore, the results of this extension confirm that an increase in dividend taxes in year t on average yields a higher propensity to pay dividends in the short run (i.e., in year t (baseline) and year t+1 (extension 2)) as the coefficient of Dividend Tax in year t+1 has a similar significance and magnitude as in year t. Also, the coefficient in column (3) is positive but, unlike in the baseline regression, highly significant suggesting that a one-pp rise in dividend taxes in year t increases the relative amount of dividends in year t+1 by 0.02pp of lagged total assets. This implies that firms pay a larger relative amount of dividends with a certain delay (i.e., in year t+1 (exten-

sion 2)) but not in the immediate aftermath of a change in dividend taxes (i.e., year t (baseline)). The positive coefficient, however, again can neither be explained by the new view nor by empirical studies evidencing a negative relation between dividend taxes and payout and thus could be an alley of future research. Contrary to the baseline results, the positive but insignificant coefficient in column (2) suggests that a change in dividend taxes on average does not change a firm's likelihood to increase or initiate dividends in year t+1. This implies that a change in dividend taxes only has an immediate impact (i.e., in year t (baseline)) on a firm's likelihood to increase or initiate dividends. Thus, the effect of Dividend Tax on Dividend Increase (t+1) is consistent with the "dividend tax neutrality" (Chetty and Saez, 2010, p.5).

Regarding personal income taxes, the results in column (1) and (2) differ from the baseline case. The coefficient in column (1) remains positive but becomes insignificant suggesting that the direct effect does not dominate the indirect effect in terms of a firm's propensity to pay dividends in year t+1; in other words, the direct effect on the variable Dividend Payer dominates the indirect effect only in the year when a change in personal income taxes occurs (i.e., year t (baseline)). Contrary to the baseline model, the indirect effect appears to prevail over the direct effect in column (2). Hence, higher labour costs incurred due to higher personal income taxes incentivise more firms to increase or initiate dividends in the year after the tax rate change (i.e., year t+1) implying that investing in corporate projects becomes increasingly unattractive compared to distributing dividends in year t+1. Only the coefficient in column (3) is similar to the baseline result: An increase in personal income taxes in year t yields a lower relative amount of dividends in year t+1 which validates the prevailing direct effect on the variable Dividend

Yield in the short run (i.e., year t (baseline) and year $t+1$ (extension 2)).

Similar to the baseline results, a rise in consumption taxes increases a firm's propensity to pay and the relative amount of dividends in year $t+1$, too, which is again in line with the predictions of the new view in period t . This implies that higher consumption taxes increase the relative attractiveness of directly distributing dividends in the short run (i.e., in year t and year $t+1$) instead of investing in corporate projects whose profit margins diminish due to the increased tax wedge. By contrast, an increase in consumption taxes on average yields no effect on a firm's likelihood to increase or initiate dividends in year $t+1$ (i.e., coefficient is insignificant but, surprisingly, negative). Thus, firms only appear to exhibit a higher likelihood to increase or initiate dividends in the immediate aftermath of a tax rate change (i.e., in year t (baseline)).

7. Robustness of Baseline Results

I test the robustness of my baseline results in two ways. First, I change the specification of my fixed effects to test whether my baseline results still hold when choosing alternative control groups. I therefore compare firms within the same country-group-industry-year where each country is clustered by (a) geographic region (i.e., countries are matched to one of the seven world regions defined by the World Bank) and (b) income group (i.e., countries are matched to one of the four income groups defined by the World Bank) instead of grouping countries by economic similarity in terms of the GDP level and GDP growth rate. Hence, I replace GDP-growth-Ln(GDP per capita)-cluster-industry-year fixed effects by (a) region-industry-year and (b) income-group-industry-year fixed effects. Second, I include the three additional country-level variables Openness, Deficit, and Interest Payments in control vector $\Gamma_{j,t}$ ³⁴. This allows me to rule out the concern that at least one of these newly included variables is a significant determinant of a firm's payout behaviour (i.e., omitted variable bias occurs) and that "nearly any desired result can be obtained" (Jacob and Jacob, 2013b, p.1259) when selecting a different set of control variables.

Table 7 presents the results of my baseline regression when including alternative fixed effects. Clustering countries by geographic regions (columns (1) to (3)) and income groups (columns (4) to (6)) mostly yields different results compared to the baseline model. All coefficients either vary in their magnitude or significance or both with the exception of dividend taxes and personal income taxes in columns (1) and (4) and columns (3) and (6), respectively. Surprisingly, the coefficient of Corporate Tax is significant implying that a rise in the corporate tax rate negatively impacts a firm's

likelihood to increase or initiate dividends (clustered by region), relative amount of dividends (clustered by region), and propensity to pay dividends in year t (clustered by income). This finding stands in stark contrast to the results of my baseline regression suggesting that corporate taxes do not affect corporate payout. Despite clustering countries by region, the coefficients of Dividend Tax are mostly similar to the baseline model. Dividend Tax, however, shows different coefficients in columns (5) and (6) if countries are grouped by income. Regarding personal income taxes, income-group-industry-year fixed effects reveal results which are mostly similar to the baseline model whereas region-industry-year fixed effects show a similar magnitude of coefficients in columns (1) and (3) but a different significance of the coefficient in column (1). Interestingly, a change in consumption taxes hardly plays a role in payout decisions when different fixed effects are employed. Even though the magnitude of coefficients in columns (3), (4), and (6) is comparable to the baseline model, they are not significant in alternative specifications. This finding stands in stark contrast to the baseline results as this robustness test suggests that consumption taxes do not affect payout.

Table 8 presents the results of my baseline regression when including additional country-level variables in control vector $\Gamma_{j,t}$. The results of this model only partially resemble the results of the baseline model. Consistent with the baseline specification, an increase in consumption taxes positively impacts all payout variables while coefficients are similarly significant with similar magnitude. Also, the effect of a change in personal income taxes on Dividend Yield (t) is in line with the baseline model due to a negative and highly significant coefficient of similar magnitude. However, the coefficient of Personal Income Tax on Dividend Payer (t) becomes insignificant and even slightly positive suggesting that personal income taxes do statistically not impact a firm's propensity to pay dividends which is not in line with my baseline results. Regarding dividend taxes, the results differ vastly from my baseline regression as no coefficient is significant at all with a negative sign across all payout variables. Even though two coefficients of Corporate Tax in columns (1) and (2) become positive, a change in corporate taxes does again not impact corporate payout which is consistent with my baseline model.

Overall, the majority of baseline results are not robust to the inclusion of different fixed effects and additional country-level controls. The only result appearing to be fully robust to alternative regression specifications is the coefficient of Personal Income Tax on Dividend Yield (t) which, in most cases, is highly significant with a similar magnitude as in the baseline model. By contrast, the effect of other tax variables on payout highly depends on the specification and thus, I derive the following two conclusions. First, the choice of the fixed effect is critical. Second, I cannot rule out the fact that my baseline model might suffer from omitted variable bias even though the low coverage of newly included variables reduces the number of observations by one quarter compared to the baseline model.

³⁴Government Debt is still excluded due to substantially poorer coverage of merely 196,656 observations compared to Openness, Deficit, and Interest Payments with a coverage of 304,225, 269,554, and 279,996 observations, respectively.

Table 7: Robustness of Main Results to Different Fixed Effects

This table shows the results of my baseline regression from 1999 to 2013 when employing different fixed effects. I replace gdp-cluster-industry-year fixed effects by region-industry-year fixed effects and income-group-industry-year fixed effects in columns (1) to (3) and columns (4) to (6), respectively. Both region and income-group follow definitions provided by the World Bank. The dependent variables are Dividend Payer (t) (column (1) and (4)), Dividend Increase (t) (column (2) and (5)), and Dividend Yield (t) (column (3) and (6)). All independent variables are defined in the appendix in Table A.1. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Corporate Tax	-0.3664 (0.2219)	-0.4609*** (0.1519)	-0.0205* (0.0120)	-0.4678** (0.1857)	0.0514 (0.2342)	-0.0164 (0.0149)
Dividend Tax	0.2127*** (0.0637)	-0.2738*** (0.0745)	0.0049 (0.0046)	0.2126*** (0.0606)	-0.0794 (0.0702)	0.0144*** (0.0040)
Personal Income Tax	-0.3613 (0.2204)	0.1851 (0.1154)	-0.0437** (0.0167)	-0.4662*** (0.1535)	0.3108 (0.2139)	-0.0499*** (0.0141)
Consumption Tax	0.7412 (0.8597)	0.9547 (0.6903)	0.0762 (0.0494)	1.2320 (0.9408)	1.4067* (0.7678)	0.0760 (0.0635)
Observations	166,133	159,769	166,133	166,131	159,770	166,131
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
GDP-Cluster-Industry-Year FE	No	No	No	No	No	No
Region-Industry-Year FE	Yes	Yes	Yes	No	No	No
Income-Group-Industry-Year FE	No	No	No	Yes	Yes	Yes
Adjusted R-squared	0.795	0.160	0.697	0.794	0.154	0.694

Table 8: Robustness of Main Results to Additional Country-level Controls

This table shows the results of my baseline regression from 1999 to 2013 when employing additional country-level variables in control vector $\Gamma_{j,t}$. I additionally include variables Openness, Deficit, and Interest Payments. Variable Government Debt is still omitted due to poor coverage. The dependent variables are Dividend Payer (t) (column (1) and (4)), Dividend Increase (t) (column (2) and (5)), and Dividend Yield (t) (column (3) and (6)). All independent variables are defined in the appendix in Table A.1. I include firm fixed effects and gdp-cluster-industry-year fixed effects in all three regressions. Please note that gdp-cluster is a placeholder representing a GDP-Growth-Ln(GDP per Capita) cluster. I report robust standard errors clustered at the country level which are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)
Corporate Tax	-0.1404 (0.1696)	-0.2452 (0.1833)	-0.0157 (0.0165)
Dividend Tax	0.2468*** (0.0654)	-0.2299** (0.1103)	0.0084 (0.0056)
Personal Income Tax	-0.2838** (0.1136)	-0.1406 (0.0982)	-0.0533*** (0.0131)
Consumption Tax	1.0857 (0.6712)	1.8513*** (0.6872)	0.0797** (0.0391)
Observations	158,184	152,337	158,184
Controls	Yes	Yes	Yes
GDP-Cluster-Industry-Year FE	Yes	Yes	Yes
Adjusted R-squared	0.799	0.164	0.703

8. Conclusion

This thesis examines the effect of corporate taxes, dividend taxes, personal income taxes, and consumption taxes on corporate payout. For this, I use a cross-country panel consisting of 115 countries over the period 1999 to 2013 and run linear regressions of the four taxes on three dependent variables measuring dividend payout. The results of the baseline regression and subsequent extensions only partially

confirm the predictions of the new view on short-run payout responses (i.e., responses in period t in the simplified single-period model). Inconsistent with initial hypotheses, corporate taxes on average do not impact dividend payout in the same year when a tax rate change becomes effective in terms of frequency and relative amounts, but a change in dividend taxes yields a statistically significant payout response even though the magnitude is small and the direction of the ef-

fect depends on the payout variable. Consistent with initial expectations, changes in personal income taxes (here: direct effect) and consumption taxes trigger mostly large payout responses. Also, cash-rich firms respond more strongly to a change in personal income taxes, consumption taxes, and (only on the country-industry level) corporate taxes. The results on payout one year after a tax rate change are mostly similar to the baseline model.

The analysis of this thesis, however, is limited to only one aspect of corporate payout (i.e., dividends) and only one part of neoclassical theory (i.e., new view in period t). In order to draw clear policy recommendations, it is therefore imperative to adopt a more holistic view by extending the scope of this thesis and investigating alternative explanations for the findings which are not in line with neoclassical theory. One way of achieving this could involve the analysis of total payout (i.e., share repurchases plus dividends; see also Chetty and Saez, 2005) since (i) share repurchases have gained importance over the last decades in the U.S. and Europe (Von Eije and Megginson, 2008) and (ii) share repurchases and dividends are, to a certain extent, interchangeable payout channels implying that a tax rate change might lead to dividends being substituted by share repurchases and vice versa (Chetty and Saez, 2005). In this context, the relative taxation of dividends vis-à-vis capital gains has to be considered, too (see also Jacob and Jacob, 2013a). Another way of deriving holistic implications involves the consideration of agency models (e.g., Chetty and Saez, 2010) which might also explain some deviations of the regression results from neoclassical predictions. Thus, I would recommend incorporating ownership structure or alternative proxies for shareholder conflicts (Jacob and Michaely, 2017) and agency issues (Chetty and Saez, 2005; Jacob and Michaely, 2017) into the regression model. Signalling models (e.g., Gordon and Dietz, 2006) could be taken into account, too, for which many executives would have to be interviewed to test whether a payout response deviating from neoclassical predictions might be interpreted as a “signal of managerial confidence in future earnings” (Jacob and Jacob, 2013a, p.188). Finally, it is possible that firms anticipated a change in taxes in previous periods such that the response in year t rather matches the predictions of the new view in period $t+1$ and is therefore a question worth being pursued in future research.

Overall, the results of this thesis should be regarded as a starting point and give managers, shareholders, and policy makers a first impression of how taxes impact corporate payout decisions which are, nonetheless, still to be complemented by future research.

References

- Alesina, A., Ardagna, S., Perotti, R., and Schiantarelli, F. Fiscal policy, profits, and investment. *American economic review*, 92(3):571–589, 2002.
- Alstadsæter, A., Jacob, M., and Michaely, R. Do dividend taxes affect corporate investment? *Journal of Public Economics*, 151:74–83, 2017.
- Auerbach, A. J. Wealth maximization and the cost of capital. *The Quarterly Journal of Economics*, 93(3):433–446, 1979.
- Auerbach, A. J. and Hassett, K. A. On the marginal source of investment funds. *Journal of Public Economics*, 87(1):205–232, 2003.
- Becker, B., Jacob, M., and Jacob, M. Payout taxes and the allocation of investment. *Journal of Financial Economics*, 107(1):1–24, 2013.
- Bradford, D. F. The incidence and allocation effects of a tax on corporate distributions. *Journal of Public Economics*, 15(1):1–22, 1981.
- Brav, A., Graham, J. R., Harvey, C. R., and Michaely, R. The effect of the may 2003 dividend tax cut on corporate dividend policy: empirical and survey evidence. *National Tax Journal*, pages 381–396, 2008.
- Chetty, R. and Saez, E. Dividend taxes and corporate behavior: Evidence from the 2003 dividend tax cut. *The Quarterly Journal of Economics*, 120(3):791–833, 2005.
- Chetty, R. and Saez, E. Dividend and corporate taxation in an agency model of the firm. *American Economic Journal: Economic Policy*, 2(3):1–31, 2010.
- DeAngelo, H., DeAngelo, L., and Stulz, R. M. Dividend policy and the earned/contributed capital mix: a test of the life-cycle theory. *Journal of Financial economics*, 81(2):227–254, 2006.
- Federal Ministry of Justice and Consumer Protection. Einkommenssteuergesetz (EStG). Retrieved on May 17, 2018, 2017. URL <https://www.gesetze-im-internet.de/estg/ESTG.pdf>.
- Feldstein, M. S. Corporate taxation and dividend behaviour. *The Review of Economic Studies*, 37(1):57–72, 1970.
- Gale, W. G. and Samwick, A. A. Effects of income tax changes on economic growth. *Working Paper*, 2016.
- Gordon, R. and Dietz, M. Dividends and taxes. Technical report, Working Paper, 2006.
- Harberger, A. C. The incidence of the corporation income tax. *Journal of Political economy*, 70(3):215–240, 1962.
- Harberger, A. C. *Efficiency effects of taxes on income from capital*. 1966.
- Jacob, M. and Jacob, M. Taxation and the cash flow sensitivity of dividends. *Economics Letters*, 118(1):186–188, 2013a.
- Jacob, M. and Jacob, M. Taxation, dividends, and share repurchases: Taking evidence global. *Journal of Financial and Quantitative Analysis*, 48(4):1241–1269, 2013b.
- Jacob, M. and Michaely, R. Taxation and dividend policy: the muting effect of agency issues and shareholder conflicts. *The Review of Financial Studies*, 30(9):3176–3222, 2017.
- Jacob, M., Michaely, R., and Müller, M. A. Consumption taxes and corporate investment. *Forthcoming in: Review of Financial Studies*, 2018.
- Jensen, M. C. Agency costs of free cash flow, corporate finance, and takeovers. *The American economic review*, 76(2):323–329, 1986.
- King, M. A. *Public policy and the corporation*, volume 3. Chapman and Hall; New York: Wiley, 1977.
- Lewellen, J. and Lewellen, K. Internal equity, taxes, and capital structure. *Unpublished Working Paper, Dartmouth College*, 2006.
- Office of the Law Revision Council. United States Internal Revenue Code. Retrieved on May 17, 2018, 2018. URL <http://uscode.house.gov/download/download.shtml>.
- Palić, I., Žmuk, B., and Grofelnik, B. The long-run impact of personal income taxation on economic development: Evidence from Croatia. *Croatian Review of Economic, Business and Social Statistics*, 3(1):35–44, 2017.
- Poterba, J. Taxation and corporate payout policy. *American Economic Review*, 94(2):171–175, 2004.
- Poterba, J. M. and Summers, L. H. New evidence that taxes affect the valuation of dividends. *The Journal of Finance*, 39(5):1397–1415, 1984.
- Poterba, J. M., Hall, R. E., and Hubbard, R. G. Tax policy and corporate saving. *Brookings Papers on Economic Activity*, 1987(2):455–515, 1987.
- Sinn, H.-W. The vanishing harberger triangle. *Journal of Public Economics*, 45(3):271–300, 1991.
- Von Eije, H. and Megginson, W. L. Dividends and share repurchases in the European Union. *Journal of Financial Economics*, 89(2):347–374, 2008.
- Wu, C. Taxes and dividend policy. *International Review of Economics & Finance*, 5(3):291–305, 1996.



Personal Taxes and Corporate Investment

Frédéric Herold

WHU - Otto Beisheim School of Management

Abstract

In this thesis, I present empirical evidence on the effect of personal taxes on firm-level investment. Exploiting a cross-country panel that consists of 40,608 firms from a total of 115 countries in the period 1999-2013, I employ a linear regression model in which I regress five different definitions of the personal tax wedge against capital investment of firms. I find that the average investment response of firms strongly depends on the definition of the personal tax wedge. My baseline regression reveals that, if the pure personal tax rate increases, firms on average show a positive capital investment response. That is, if firms cannot shift the economic burden of personal taxes to other stakeholders, an increase in personal taxes, *ceteris paribus*, increases the factor price of labour and thus exerts higher pressure on corporate profits. Profit-maximising firms therefore counteract this pressure by (partially) substituting the more expensive input factor labour by capital, increasing their capital investment. This effect, however, does not hold true for alternative definitions of the personal tax wedge that additionally include social security contributions. Likewise, I obtain mixed results when testing for cross-sectional variation in capital investment responses arising from differences in relative market power, the ability to substitute input factors, and financial constraints. In this context, my thesis provides empirical evidence on the effect of personal taxes on investment behaviour at the firm level and thus adds to current literature, which mainly considers the effect of personal taxes on aggregate investment, economic growth, and total factor productivity.

Keywords: investment; personal tax; tax wedge

1. Introduction

Over the past decades, a substantial amount of literature has evolved which extensively discusses the effect of corporate taxes (e.g., Auerbach et al., 1983; Djankov et al., 2010; Dobbins and Jacob, 2016; Giroud and Rauh, 2017; Ljungqvist and Smolyansky, 2016), payout taxes (e.g., Alstadsæter et al., 2017; Becker et al., 2013; Chetty and Saez, 2010; Yagan, 2015), and consumption taxes (e.g., Jacob

et al., 2018) on investment behaviour of firms¹. The discussion on the effect of personal taxes on firm-level investment, however, is much more fragmentary and less diverse. That is, although previous literature on personal taxes does exist, evidence on the direct effect of personal taxes on firm-level investment is surprisingly scarce. For instance, one set of studies exclusively relies on macroeconomic data and draws unclear conclusions about the effect on aggregate invest-

I would like to express my very great appreciation to Prof. Dr. Martin Jacob, my first supervisor, for his extensive guidance and structured advice during this research project. In particular, I would like to thank him for proposing the topic "Personal Taxes and Corporate Investment", providing a pre-edited sample of core data, and engaging in fruitful discussions with me that altogether helped me successfully write this thesis. I would also like to express my deep gratitude to Dr. Antonio De Vito, my second supervisor, for his far-reaching support through insightful deep-dive sessions on specialised statistical software and meaningful intellectual discussions that shaped the outcome of this research project. Finally, I would like to thank Nicholas Herold, Ralf Herold, and Dr. Tobias Wagner for their helpful comments and suggestions that improved the quality of this thesis.

¹Dobbins and Jacob (2016) provide a comprehensive overview of studies which discuss a negative effect of corporate taxes on investment, both for the macro level (e.g., Auerbach et al., 1983; Djankov et al., 2010) and the direct effect on firm-level investment (e.g., Dobbins and Jacob, 2016; Ljungqvist and Smolyansky, 2016). Similarly, the effect of payout taxes on investment levels in the light of agency issues (e.g., Alstadsæter et al., 2017; Chetty and Saez, 2010) and the allocation of investment between cash-rich and cash-poor firms (e.g., Alstadsæter et al., 2017; Becker et al., 2013; Yagan, 2015), although with mixed empirical results, has been extensively investigated. Also, Jacob et al. (2018) provide recent empirical evidence on the effect of consumption taxes on firm-level investment which complements previously inconclusive findings on the macroeconomic level (e.g., Alesina et al., 2002; Arnold et al., 2011).

ment and economic growth² (e.g., Lee and Gordon, 2005). Other studies, by contrast, attempt to complement these macro-level findings by estimating the effect on total factor productivity (e.g., Arnold et al., 2011) or by employing the q approach (e.g., Alesina et al., 2002) but they show no statistically significant, robust effect on firm-level investment³. Thus, it appears that previous studies have unclear implications for investment responses on the firm level which creates a substantial gap in tax research.

This neglect is astonishing when considering the importance of personal taxes for fiscal budgets and their practical relevance for input factor decisions of firms. First, personal taxes are a major source of tax revenues on the fiscal level and on average contribute to approximately 25% of tax revenues in OECD countries (Organisation for Economic Co-operation and Development, 2017) which emphasises the significance of personal taxes as a policy instrument. Second, and even more severely, if firms cannot fully pass the economic burden of personal taxes onto other parties (e.g., Dyreng et al., 2017; Jacob et al., 2018), personal taxes can, ceteris paribus, distort input factor decisions on the firm level, and thus the optimal factor mix of firms by increasing the factor price of labour. When abstracting from productivity differences between factors, this ‘price increase’ is expected to reduce the attractiveness of the input factor labour in favour of capital, and thus likely creates pressure to substitute the more expensive input factor labour by additional capital. Considering these substantial implications, it is imperative for policy makers and managers to understand the effect of personal taxes on investment behaviour of firms.

This thesis therefore aims at bridging this gap by providing empirical evidence on the effect of personal taxes on firm-level investment and the magnitude of this effect. For this, my empirical analyses exploit a cross-country panel of non-financial, non-transportation, non-telecommunication, non-utility firms in 115 countries over the 1999–2013 period. My estimation strategy is threefold. First, following Jacob et al. (2018), I employ linear probability models to identify “country-level determinants of ... [personal] tax changes” (Jacob et al., 2018, p.15). Second, my baseline model in which I account for “observable firm and [country-level] characteristics” (Alstadsæter et al., 2017, p.75) and include firm- and deficit-interest-payment-cluster-industry-year fixed effects estimates the average investment response. Third, I test for cross-sectional variation in investment responses to analyse the impact of differences in firm characteristics such as market power, the ability to substitute input factors, and financial constraints on the responsiveness of capital invest-

ment. In all tests, five different definitions of the personal tax rate (i.e., one pure personal tax rate and four different specifications including social security contributions) are employed to investigate whether investment responses of firms differ depending on the definition of the personal tax wedge.

Interestingly, my empirical results reveal exactly that. In my baseline regression, for instance, I can only validate a positive average response of capital investment for the pure personal tax rate (although the effect size is smaller than for other taxes) whereas specifications including social security contributions are statistically insignificant. This finding supports my proposed mechanism of firms facing higher pressure to substitute labour by capital but does not confirm predictions about social security contributions having the same economic effect on factor decisions as the pure tax rate. This picture slightly changes when testing for cross-sectional variation in investment responses where results are partially ambiguous. For instance, if firms have low market power, investment reacts more strongly compared to the average investment response in case of the pure personal tax rate, but the response mostly reverses (i.e., investment reacts less strongly) when including social security contributions in the personal tax wedge. Results also appear to be mixed when testing for differences in the ability to substitute labour by capital and financial constraints. Hence, my thesis contributes to the literature by providing empirical evidence on the direct relationship between the personal tax rate and investment behaviour at the firm level, and thus illustrates the impact of policy instruments on input factor decisions and the optimal factor mix of firms.

The remaining sections of this thesis are structured in the following way. In section 2, the theoretical background is explained based on which I derive four hypotheses (i.e., one predicting the average investment response and three investigating cross-sectional variation in capital investment responses). Section 3 presents my data, methodology, and summary statistics on variables used in my baseline regression. Furthermore, I conduct a pre-analysis and check for sufficient variation in personal tax changes in section 4 on which I base my baseline regression and subsequent analyses of cross-sectional variation in section 5. I then test for robustness of my baseline results in section 6. Finally, my conclusion is presented in section 7.

2. Theoretical Background: Model and Hypothesis Development

2.1. Optimal Input Factor Mix and Personal Tax Wedge

According to economic theory, the “production function [of firms] has two input factors, capital and labor” (Dobbins and Jacob, 2016, p.8). However, since firms are an investment vehicle of their shareholders (Alstadsæter and Jacob, 2012), and thus are assumed to be profit-maximising entities, they must decide on the optimal mix of these factors to produce a certain output at minimal costs. Following Pindyck and Rubinfeld (2018), the optimal factor mix is determined

²Lee and Gordon (2005) admit that “the aggregate information reported ... is insufficient to draw ... conclusion[s] about ... links between [personal] tax[es] ... and growth” (p.15).

³Arnold et al. (2011) investigate the effect of personal taxes on industry-level entrepreneurial activity and total factor productivity but fail to do so for firm-level investment. Likewise, Alesina et al. (2002) “estimate a q type of investment equation that links investment to ... profits” (p.572) but they solely rely on aggregate measures such as “investment of the business sector ... [and] capital stock” (p.578).




Figure 1: Personal Tax Wedge




Figure 2: Substitution Response After a Personal Tax Increase

by the two criteria (a) factor productivity⁴ and (b) price per input factor unit. That is, the more output a factor can produce within a certain time (i.e., the more productive a factor) for given factor price, the higher its contribution for the generation of revenues, and thus the more attractive the input factor. Likewise, the lower the price of a factor for a given productivity level, the higher the profit margin per unit of output produced, and hence the more attractive the input factor. Thus, when combining these two criteria, the optimal factor mix is a function of the relative attractiveness of input factors which can be expressed as the ratio of factor productivity to factor price⁵.

Personal taxes, however, can change the optimal factor mix of firms. As illustrated in Figure 1, the “tax wedge theory” (Becker et al., 2013, p.5; see also Alstadsæter et al., 2017; Jacob et al., 2018) predicts that personal taxes drive a wedge between the factor price of labour paid by firms (c_L) and the net wage of employees (w_n). Thus, unless firms can fully shift “the economic burden, or incidence, [of personal taxes]” (Dyregang et al., 2017, p.6) to consumers via

⁴Factor productivity is defined as the level of output which can be produced by an input factor within a given time.

⁵For simplicity, I assume that the relative attractiveness of input factors only changes the mix of input factors whereas the level of output generated remains constant irrespective of the input factor mix. I also abstracted from other determinants of factor decisions, e.g., the availability of input factors (which is assumed to be reflected in the price) and the state of technology.

higher market prices or workers via lower net wages (Dyregang et al., 2017), personal taxes increase the factor price of labour while labour productivity remains constant⁶, and thus they reduce the attractiveness of labour relative to capital.

Consequently, personal taxes exert pressure on profits, and thus force profit-maximising firms to substitute the relatively more expensive factor labour by additional capital⁷. Figure 2 visualises this relationship by using a simplified P&L structure which assumes firms to bear part of the personal tax incidence.

To conclude, personal taxes are expected to discriminate the input factor labour in favour of the input factor capital, and thus distort input factor decisions of firms⁸. Based on this, I develop four hypotheses on the investment behaviour

⁶I expect the higher factor price of labour not to be offset by increases in labour productivity (although this could be assumed in a world without personal taxes in which employees are paid a wage equal to their marginal productivity (Pindyck and Rubinfeld, 2018)). Thus, ceteris paribus, a tax-induced increase in the factor price of labour results in a lower attractiveness of labour relative to capital.

⁷I assume that labour and capital are, on the margin, substitutes (e.g., Dyregang et al., 2017; Jacob et al., 2018). Please refer to hypothesis one in section 2.2 for a detailed explanation. For a substitution response to be economically reasonable, capital is also assumed to have a productivity greater than zero, and firms are assumed to keep their output level constant.

⁸In a wider sense, personal taxes can be a variable not just including the top marginal income tax rate on labour income, τ^P , but also labour-related costs such as social security contributions which drive a wedge between w_g

of firms. In hypothesis one, I predict the average investment response. Hypotheses two, three, and four, then extend the scope of my model and capture cross-sectional variation in the responsiveness of capital investment.

2.2. Hypothesis Development

Hypothesis 1: On average, if the economic burden of a personal tax increase is (partially) borne by firms, capital investment responds ambiguously.

Assuming supply and demand to be neither fully elastic nor inelastic (e.g., Jacob et al., 2018) in the labour market, the economic burden of a personal tax increase is shared between firms and employees (i.e., higher labour costs for firms, lower net wage for employees). At the firm level, this exerts higher pressure on profits, and thus forces profit-maximising firms to reduce costs incurred by their deployment of input factors. That is, since an increase in personal taxes directly increases the factor price of labour, firms would unambiguously try to reduce their labour intake in their production function to cut costs.

The effect on capital investment, however, is ambiguous and depends on whether labour and capital, on the margin (i.e., in marginal factor decisions), are complements or substitutes. Two channels of investment responses are hence plausible (e.g., Dyreng et al., 2017). First, like for labour, firms can respond by reducing their capital investment, too. This would allow them to “maintain their [optimal] mix of input factors” (Dobbins and Jacob, 2016, p.4), for which labour and capital, even on the margin, would be treated as complements. Second, by contrast, capital investment of firms could increase. Such a response would occur if labour and capital could be partially (i.e., to a small extent) substituted despite their overall complementarity, and thus both input factors would be substitutes on the margin. The second channel is empirically supported by Dyreng et al. (2017) showing that labour and capital, on the margin, can be substitutes.

Hypothesis 2: After an increase in personal taxes, firms with low market power vis-à-vis their stakeholders show greater responsiveness in capital investment.

Intuitively, the personal tax incidence borne by firms (and ultimately shareholders) likely determines the magnitude of investment responses. That is, the greater (smaller) the economic burden on firms, the greater (smaller) the pressure to substitute labour by capital. Yet, previous literature suggests that “shareholders might not bear the entire economic burden [of personal taxes] ... [since a] firm’s market power

and w_n . These additional labour costs are effectively part of the gross wage, wg_n , and thus are expected to have the same economic effect on firm-level investment as the pure personal tax rate τ^p . Although these labour-related costs are no taxes, social security contributions will nevertheless be included in the definition of the personal tax wedge in section 5 to check whether they empirically have the same economic effect on investment.

allows it to pass the [economic] burden to [stakeholders such as suppliers,] workers, or consumers” (Dyreng et al., 2017, p.1), and thus cross-sectional variation in investment responses may result from differences in the relative market power of firms. Since market power is a function of market demand elasticity (in the case of consumers) and supply elasticity (in the case of suppliers/workers) (e.g., Dyreng et al., 2017; Jacob et al., 2018), I present two cases in a partial equilibrium setting which show the relationship between market power and firm-level investment⁹. Also, the model of the profit-maximising firm is assumed (e.g., Dyreng et al., 2017) that will try to reduce the economic burden imposed by personal taxes.

Conceptually, it does not matter onto which stakeholder the economic burden of a personal tax increase is shifted as investment responses of firms are unambiguous in both subsequent cases. First, I consider the market power of firms vis-à-vis their employees on the cost side¹⁰. In this case, market power depends on the elasticity of labour supply (e.g., Dyreng et al., 2017; Jacob et al., 2018). That is, the more elastic (inelastic) the labour supply (e.g., due to high (low) education levels (e.g., Dyreng et al., 2017; Fuest et al., 2018) and correspondingly high (low) labour mobility), the lower (higher) the ability of firms to freely set wages, and thus the lower (higher) their ability to shift the economic burden of a personal tax increase to employees¹¹. Subsequently, this exerts higher (lower) pressure on profits, and thus increases the (creates less) pressure to substitute the more expensive factor labour by capital which, in turn, causes investment of firms with low (high) market power to respond more (less) strongly¹² than the average investment response. Second, I consider the market power of firms vis-à-vis their consumers which is a function of the elasticity of market demand on their revenue side (e.g., Dyreng et al., 2017; Jacob et al., 2018). That is, the more (less) elastic the market demand (e.g., due to the availability of substitutes (Jacob et al., 2018), the lower (higher) the ability of firms to shift the economic burden of a personal tax increase to consumers through higher prices. Thus, this translates into higher (creates less) pressure to substitute labour by capital and is expected to result in a stronger (weaker) investment response if firms have low (high) market power.

⁹For illustration purposes, I abstract from a general equilibrium setting in which “firm[s] can simultaneously shift [their personal] tax burden to [multiple stakeholders]” (Dyreng et al., 2017, p.10).

¹⁰Literature suggests that “results are essentially the same [if] firms ... pass on taxes to ... suppliers through [lower] input ... prices instead of passing [them] on ... to workers [through lower wages]” (Jacob et al., 2018, p.2).

¹¹Alternatively, it could be argued that the power of unions influences the ability of firms to shift the economic burden to employees. However, union power belongs to the discipline of bargaining literature (e.g., Katz, 1993) from which I abstract in my model for simplicity.

¹²The meaning of more strongly depends on the direction of the average effect. That is, if the average effect is positive (negative), I expect a stronger increase (decrease) in investment if firms have low market power.

Hypothesis 3: After an increase in personal taxes, capital investment responds more strongly if firms can more easily substitute labour by capital.

As illustrated in Figure 2, an increase in personal taxes exerts stronger pressure on firms to substitute the more expensive input factor labour by capital¹³. This does, however, not imply that firms are able to substitute both factors to the same extent, and thus cross-sectional variation in investment responses across firms may arise from differences in the ability to substitute labour by capital (e.g., *Dyreg et al., 2017*). Intuitively, the degree of input factor substitutability is influenced by two elements: (a) The knowledge-intensity of the factor labour, and (b) the importance of labour and capital for the generation of value added (i.e., output produced).

First, knowledge-intensive labour tasks such as R&D are difficult to automate, and thus the more knowledge-intensive the factor labour, the more difficult it is to substitute labour by capital. Consequently, firms for which knowledge and innovation (i.e., R&D) are core of their business model (e.g., consultancies or pharma firms) have a high share of knowledge-intensive labour, and thus have a lower ability to substitute labour by capital, if they can do so at all¹⁴. Second, the higher the importance of an input factor in a firm's production function, and hence for the generation of value added (i.e., output produced), the more difficult it is to substitute this input factor if output is to be kept constant. For example, if labour (capital) is highly productive and therefore important for the generation of output, the same output can, if at all, only be produced by a disproportionately high amount capital (labour), and hence it is relatively more difficult (easier) to be substituted by labour by capital on the margin. Thus, the greater a firm's ability to substitute labour by capital, the more strongly investment is expected to respond since firms likely show a smaller (greater) substitution response towards capital if labour is knowledge intensive (capital is important for the generation of output).

Hypothesis 4: After an increase in personal taxes, financially constrained firms which strongly rely on internal cash flows for investment exhibit a more negative investment response.

Besides differences in relative market power and the ability to substitute input factors, previous literature suggests that cross-sectional variation in investment responses can also result from "differences in the availability of internal funds" (*Jacob et al., 2018, p.5*) across firms. That is, if internal cash flows are the marginal source of finance, investments in cash-constrained firms (i.e., firms with limited

internal resources) are likely more prone to decreases in internal cash flows than investments in cash-rich firms (i.e., firms with abundant internal resources) (e.g., *Dobbins and Jacob, 2016; Faulkender and Petersen, 2012; Fazzari et al., 1988; Jacob et al., 2018*). Thus, assuming firms to bear part of the personal tax incidence, an increase in personal taxes is predicted to reduce profits, and thus internal after-tax cash flows and investments more strongly if firms are cash-constrained and "heavily [rely] on internally generated cash flows for investment" (*Jacob et al., 2018, p.5*). Thus, despite higher pressure to substitute labour by capital, this effect is expected to translate into a more negative investment response of financially constrained firms¹⁵ as their availability of internal resources is more strongly affected.

3. Data, Methodology, and Summary Statistics

The data used in this thesis have largely been provided by the WHU chair of Business Taxation and stem from four main data sources. First, firm-level data on listed companies over the 1997–2013 period were retrieved from the Compustat Annual North America and Global database. Second, tax policy data were extracted from handbooks published by major auditing and tax advising firms such as KPMG, PwC, Ernst & Young, and Deloitte and are available from 1999–2013. Third, information on macroeconomic and governance indicators follow the World Bank definition and originate from the World Bank website for all countries in the dataset. Fourth and finally, I retrieved additional data on personal taxes from the OECD tax database from 2000–2013 to include social security contributions in the definition of the personal tax wedge.

Prior to merging datasets, I amended the data in several ways to increase the coverage of some variables. For instance, I added new data on Tobin's q with higher coverage across firms which were provided by the WHU chair of Business Taxation¹⁶. Similarly, I replaced missing data entries of the variable Income Group to increase the number of observations for the income-group-cluster used in robustness tests of my thesis¹⁷. In addition to the datasets provided by

¹⁵This corresponds to lower investment levels of financially constrained firms compared to the average investment response, irrespective of the direction (i.e., coefficient) of the average effect. Since the average effect in hypothesis one is expected to be ambiguous (i.e., both $\beta_1 < 0$ and $\beta_1 > 0$ are plausible), the investment response of financially-constrained firms is therefore predicted to be more negative (and not greater or smaller than the average effect as such a statement requires a clear prediction of the direction of the average effect).

¹⁶The definition of Tobin's q is the same as in *Jacob et al. (2018)* (i.e., the market value of equity over total assets). It was necessary to add new data on Tobin's q since the variable Market Value (denoted by *mkvalt*) in the provided Compustat data suffered from poor coverage. Attempts to estimate this variable via share price * number of shares as in the originally provided Compustat dataset only increased the coverage marginally.

¹⁷This adjustment was carried out in two steps. First, I manually replaced missing values for Argentina, Jamaica, New Zealand, and Nigeria based on World Bank data. Second, I merged new data from the World Bank website for all other 66 countries with missing data entries to the #3.1_full_codes.dta dataset. Missing countries, for instance, included Taiwan, Cyprus, Monaco, and Paraguay.

¹³This assumes labour and capital to be substitutes on the margin.

¹⁴In this hypothesis, I abstract from recent technological developments in the field of artificial intelligence. These developments potentially increase the ability of firms to automate knowledge-intensive labour since they increasingly enable the factor capital to perform knowledge-intensive tasks (e.g., in R&D). Thus, knowledge-intensive labour could be more easily automated (and substituted by capital) in future.

Table 1: Summary Statistics of Main Variables

This table displays summary statistics of all main variables from 1997 to 2013. Panel A presents descriptive statistics for variables on the country level. Panel B summarises descriptive statistics for variables on the firm level. Please refer to Table A.1 in the appendix for variable definitions. Notes: Summary Statistics of all firm-level variables in Panel B correspond to the winsorised version of the respective variable to eliminate the effect of outliers on my results.

Variable	N	Mean	Standard Deviation	25th percentile	Median	75th percentile
Panel A: Country-level Variables						
Tax Policy Variables						
Personal Tax	345,333	0.3972	0.0938	0.3500	0.4000	0.4641
Corporate Tax	345,954	0.3215	0.0742	0.2700	0.3300	0.3900
Consumption Tax	325,864	0.1073	0.0627	0.0519	0.1000	0.1700
Payout Tax	345,333	0.1621	0.1027	0.1000	0.1500	0.2488
Accelerated Depreciation	345,954	0.8109	0.3916	1	1	1
LCB	345,954	0.4334	0.4955	0	0	1
Group Taxation	343,328	0.5521	0.4973	0	1	1
Progressive	345,954	0.6302	0.4828	0	1	1
Extended Tax Definitions						
67% Earner	201,247	0.3685	0.1017	0.3198	0.3439	0.3939
100% Earner	201,247	0.4018	0.0936	0.3423	0.3883	0.4361
133% Earner	201,247	0.4420	0.0918	0.4093	0.4336	0.4770
167% Earner	201,247	0.4275	0.0984	0.3525	0.4340	0.4748
Macroeconomic Variables						
GDP Growth	363,902	3.5813	3.4689	1.7292	3.1400	5.1472
Ln(GDP per Capita)	363,817	9.6124	1.3841	8.6600	10.4290	10.5557
Inflation	363,902	2.7073	4.3191	0.8477	2.0327	3.7157
Deficit	269,504	-2.6679	3.9788	-4.8523	-3.1779	0.0177
Openness	304,174	0.7266	0.8649	0.2829	0.4831	0.6549
Interest Payments	279,947	0.0225	0.0123	0.0150	0.0230	0.0276
Government Debt	196,624	60.9360	37.7089	40.0881	53.5029	64.0318
Governance Indicators						
Voice and Accountability	371,022	0.6717	0.8952	0.3900	1.0100	1.3500
Political Stability	371,017	0.3316	0.8166	-0.2000	0.6000	0.9600
Government Effectiveness	371,006	1.1319	0.7665	0.4000	1.4600	1.7500
Regulatory Quality	371,006	0.9722	0.7837	0.4200	1.1900	1.6200
Rule of Law	371,022	0.9888	0.7968	0.2900	1.3300	1.6100
Control of Corruption	371,006	0.9794	0.9781	0.0500	1.2900	1.8350
Panel B: Firm-level Variables						
Investment	321,987	0.0719	0.1096	0.0139	0.0357	0.0803
Cash Holdings	338,232	0.1269	0.2727	0.0020	0.0203	0.1129
Profit	337,817	0.0268	0.2106	-0.0036	0.0517	0.1141
Leverage	369,749	0.0933	0.1563	0.0007	0.0112	0.1167
Ln(Sales Growth)	323,754	0.0876	0.4335	-0.0460	0.0730	0.2182
Sales Growth	287,128	0.4841	1.7148	-0.0637	0.1468	0.4652
Loss	370,210	0.2984	0.4576	0	0	1
Tobin's q	279,446	1.4997	3.5907	0.3319	0.6837	1.4016
Size	388,193	6.5558	3.0442	4.3872	6.3843	8.4709

the WHU chair, I retrieved and added data on geographic regions following World Bank definition from the World Bank website to construct a region-cluster later in my robustness

section.

After merging datasets, I conducted general data cleaning to eliminate implausible observations. For instance, I

dropped firms with SIC codes 4000 to 4999 (i.e., utility, transportation, and telecommunication firms) and 6000 to 6999 (i.e., financial firms) since these subsets of firms likely exhibit different investment behaviour which could distort my results¹⁸. In addition, I excluded observations for which firms had negative total assets or for which total assets were unavailable. Likewise, I limited my baseline sample to observations with positive sales and cash holdings¹⁹. To eliminate bankrupt firms, I also dropped observations with a value of Common/Ordinary Equity smaller than or equal to zero and which possess a Leverage ratio greater than or equal to one. The sample was further limited to observations for which the macroeconomic variables GDP per capita, Openness, Government Debt, and Interest Payments were not negative to eliminate further implausible observations. Additionally, I conducted specific data cleanings tailored to my research question. For instance, I only included observations for which capital expenditure was greater than or equal to zero to restrict my analyses to firms with non-negative investment. Similarly, I dropped implausible tax rates with values less than zero or higher than one. I also conducted specific data cleanings in my cross-sectional variation analyses (e.g., by dropping negative (i.e., implausible) net PPE when testing for different factor substitutability across firms), but these cleanings were carried out after my baseline tests and thus do not affect the composition of my baseline sample or robustness tests. Following my data cleanings, I converted firm-level variables which were denoted in currencies other than U.S. Dollar into U.S. Dollar by using the average annual U.S. Dollar exchange rate in the corresponding year issued by the WHU Chair of Business Taxation²⁰. In addition, I winsorised all non-dummy, firm-level variables and their lags below the 1st and above the 99th percentile to reduce the effect of extreme outliers on my results²¹. Overall, these adjustments result in a baseline sample comprising 42,670 firms located in 115 countries from 1997–2013²². Table 1 presents summary statistics on all variables used in my baseline specifica-

tion after these adjustments.

4. Pre-Analysis: Linear Probability Model and Variation in Personal Tax Rate Changes

Prior to running regressions on corporate investment behaviour, my data on personal tax rates must fulfil two fundamental conditions. First, my independent variable of interest (i.e., the personal tax rate) must exhibit a sufficiently large degree of variation in my sample. Otherwise, my causal inference would be limited to a few selected events and could barely be generalised to all countries available in my dataset (Jacob et al., 2018). Fortunately, my cross-country panel of 115 countries provides a solid source of tax rate variation as personal taxes change 217 times from 1999 to 2013 (thereof 76 increases and 141 decreases). Even when abstracting from personal tax changes of less than two percentage points, 121 changes can still be observed (thereof 43 increases and 78 decreases). Consequently, my dataset shows a sufficiently large variation of the personal tax rate and fulfils the first condition.

Second, changes in the personal tax rate must be exogenous to allow for clear causal inference. This is especially critical since my baseline regression assumes changes in the personal tax rate to be entirely exogenous. Otherwise, I would only “observe a spurious correlation” (Jacob et al., 2018, p.15) instead of a causal relationship between personal taxes and investment. Analogously to Jacob et al. (2018), I therefore address endogeneity concerns by running a linear probability model showing whether changes in the personal tax rate are related to the business cycle or other economic conditions. In the model, I include the six macroeconomic determinants GDP Growth, Ln(GDP per capita), Inflation, Deficit, Openness, and Interest Payments on government debt as regressors²³ (Jacob et al., 2018). Likewise, I also use country fixed effects and region-year fixed effects to capture time invariant effects at the country level and limit comparable countries to their counterparts within the same World Bank region (Jacob et al., 2018).

Table 2 displays results of my linear probability model. In columns (1) and (2), I model whether macroeconomic determinants affect the probability of personal tax changes by more than 2.0 percentage points. As the dependent variable, I use a dummy equal to one if personal taxes are increased (column 1) or decreased (column 2). In addition,

¹⁸Asker et al. (2011), for instance, argue that financial firms and utility firms are subject to different regulation affecting their investment policy. Similarly, companies in the transportation and telecommunication sector mostly tend to be formerly state-owned and, due to their business model, I expect them to possess a substantial amount of fixed assets with correspondingly high capital expenditure. It is therefore plausible to assume that these subsets of firms differ substantially in their investment behaviour compared to all other firms included in the sample (and thus could distort my results).

¹⁹Please note that cash holdings are defined as the sum of cash holdings and short-term investments because short-term investments are assumed to be as liquid as cash. Please refer to Table A.1 in the appendix for exact variable definitions.

²⁰Some firm-level variables such as EBIT, sales, or total assets were already denoted in USD. Therefore, I excluded these variables from the currency conversion process.

²¹I refrained from winsorising my tax policy variables and country-level data from the World Bank since these are official statistics. Similarly, the appended data on Tobin's q were already winsorised and hence excluded from the winsorisation process.

²²Since data on tax policy variables are only available from 1999–2013, the sample is ultimately restricted to 40,608 firms from 1999–2013 in subsequent regressions. If social security contributions are included in the defini-

tion of the personal tax wedge, the sample further shrinks to 25,874 firms as data on social security contributions are only available for OECD countries from 2000–2013.

²³In the excel file 2. LPM Results Edited.xls, three specifications of this model were used. In specification (1), I additionally included Government Debt as a regressor but abstracted from it in specifications (2) and (3). Also, specifications (1) and (3) are restricted to the same 410 observations, whereas specification (2) considers 743 observations. I therefore reported specification (2) to avoid distorted results due to a poor coverage of Government Debt. This is supported by similar results (both magnitude and significance) in specifications (1) and (3) indicating that omitting Government Debt is unlikely to cause an omitted variable bias.

Table 2: Linear Probability Model Results

This table presents results of my linear probability model. In columns (1) and (2), I model macroeconomic determinants which affect the probability of changes in the personal tax rate by more than 2.0 percentage points. The magnitude of these changes are modelled in columns (3) to (5). Please refer to Table A.1 in the appendix for definitions of explanatory variables. I further include country fixed effects and region-year fixed effects in all specifications. This table also reports robust standard errors clustered at the country level in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Probability of		Tax Change (3)	Magnitude of	
	Tax Increase (1)	Tax Decrease (2)		Tax Increase (4)	Tax Decrease (5)
GDP Growth	-0.0025 (0.0031)	-0.0054 (0.0058)	0.0000 (0.0006)	-0.0014 (0.0031)	-0.0116 (0.0070)
Ln(GDP per Capita)	-0.0948 (0.1236)	0.1222 (0.2334)	0.0114 (0.0218)	0.1483 (0.1906)	0.2556 (0.2674)
Inflation	0.0010 (0.0014)	0.0038 (0.0023)	0.0000 (0.0002)	0.0007 (0.0013)	0.0015 (0.0027)
Deficit	-0.0015 (0.0037)	0.0092 (0.0062)	-0.0005 (0.0004)	-0.0130* (0.0068)	0.0072 (0.0078)
Openness	0.0879 (0.1100)	-0.1711 (0.1574)	-0.0024 (0.0187)	0.0332 (0.1406)	0.0022 (0.1820)
Interest Payments	1.0763 (1.6420)	0.9006 (1.9015)	0.1096 (0.1674)	5.0101* (2.9223)	-0.5837 (3.4237)
Observations	743	743	743	743	743
Country FE	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	-0.012	-0.048	-0.106	0.134	0.073

the magnitude of all 217 personal tax changes is modelled in the remaining columns. In column (3), the dependent variable is denoted by the change in the personal tax rate. In columns (4) and (5), I interact this change with a dummy for a tax increase and tax decrease, respectively (e.g., [Jacob et al., 2018](#)).

Overall, based on my dataset, changes in the personal tax rate appear to be mostly exogenous since four macroeconomic variables are not significant. In addition, the probability and the magnitude of personal tax changes seem to be mostly unaffected by economic conditions except for the magnitude of personal tax increases. This is indicated by significant coefficients for Deficit and Interest Payments in column (4). That is, if the budget deficit increases (e.g., in recessions), policy makers tend to increase personal taxes less strongly, thereby limiting the adverse effect of personal taxes on economic growth. Furthermore, policy makers tend to increase personal taxes more strongly to finance higher interest payments which, for example, could be a result of formerly high budget deficits. Considering these results, I define quartiles of Deficit and Interest Payments for each year and create a deficit-interest-payment-cluster-industry-year fixed effect for my baseline regression. This assures that firms in countries with personal tax changes are compared to a control group which is subject to similar economic conditions in terms of budget deficit and interest payments.

5. Main Empirical Analysis and Results

In this section, I estimate the causal effect of a change in the personal tax rate on investment at the firm level. To accomplish this, I structured this section into two main parts. First, the average effect on investment is analysed in my baseline model using the cross-country panel of 115 countries from 1999 to 2013 (2000 to 2013 for social security contributions). Second, I examine cross-sectional variation in investment responses due to cross-sectional differences in firm characteristics such as (a) market power vis-à-vis stakeholders, (b) different degrees of input factor substitutability, and (c) the presence of financial constraints.

5.1. Baseline Regression

To estimate the average effect of personal taxes on corporate investment behaviour, I construct the following linear regression model based on the estimation method of ordinary least squares:

$$\text{Inv}_{i,j,t} = \alpha_0 + \beta_1 \text{Personal Tax}_{j,t} + \delta_1 \Gamma_{j,t} + \delta_2 T_{j,t} + \delta_3 \Phi_{i,j,t-1} + \alpha_i + \alpha_{g,k,t} + \epsilon_{i,j,t} \quad (1)$$

My dependent variable is Investment of firm i located in country j in year t . Consistent with previous literature (e.g., [Jacob et al., 2018](#)), I approximate my dependent variable with capital expenditure over lagged total assets. My independent variable of interest is the personal tax rate which is denoted by $\text{Personal Tax}_{j,t}$. I employ five different definitions

of the personal tax rate. First, the top marginal income tax rate on labour income is used to analyse the effect of the pure personal tax rate on investment. Second, I extend this definition and include social security contributions. Doing so, I consider four different income classes of employees in OECD countries, which are expressed as a percentage of the average wage earned in a respective country-year²⁴.

To account for variables which could affect investment other than personal taxes, I include three control vectors in my baseline regression. First, I account for country-level factors in vector $\Gamma_{j,t}$ which comprises the macroeconomic variables GDP Growth, Ln(GDP per Capita), Inflation, Deficit, and Interest Payments as well as the governance indicators Voice and Accountability, Political Stability, Government Effectiveness, Regulatory Quality, Rule of Law, and Control of Corruption as defined by the World Bank (e.g., Jacob et al., 2018). Variables with poor coverage such as Openness or Government Debt, however, were excluded to increase the number of observations in my regression²⁵. Second, my control vector $T_{j,t}$ contains a set of tax policy variables including Accelerated Depreciation, LCB, Group Taxation, and Progressive²⁶ analogous to Jacob et al. (2018). To address concerns that changes in the personal tax rate coincide with changes in other tax rates, I additionally include other tax rates such as Consumption Tax_{j,t}, Payout Tax_{j,t}, and Corporate Tax_{j,t} in the tax policy variable vector $T_{j,t}$ (e.g., Jacob et al., 2018). Doing so enables me to isolate the effect of personal tax changes on firm-level investment. Third and finally, I include control variables on the firm level via vector $\Phi_{i,j,t-1}$. In this vector, I account for Cash Holdings, Profit, Leverage, Ln(Sales Growth), Tobin's q, Size, and Loss analogous to previous investment literature²⁷ (e.g., Baker et al., 2003; Cummins et al., 1996; Dobbins and Jacob, 2016; Jacob et al., 2018). All firm-level controls are lagged by one period to eliminate concerns about endogeneity (Dobbins and Jacob, 2016).

Furthermore, my baseline model includes two fixed

²⁴These alternative definitions follow the definition of the OECD tax database and are conceptually no taxes. However, I nevertheless expect social security contributions to have the same economic effect on investment as the pure personal tax rate.

²⁵Please refer to Table 1 in section 3 for an overview of the coverage of main variables. My baseline results are robust to including Openness as an additional control variable when using the deficit-interest-payment-cluster-industry-year fixed effect of my baseline specification. Please refer to the excel file 3. Baseline Results Edited.xls for detailed results.

²⁶Dreßler and Overesch (2013), for instance, discuss that LCB and Group Taxation influence investment behaviour of firms. Besides, I expect Accelerated Depreciation and Progressive to affect investment decisions and risk-taking of firms, respectively. A dummy for loss carry forwards has not been included in my model as all countries allow for loss carry forwards in the sample period.

²⁷This set of firm-level controls is included for several reasons. Cash Holdings and Profit are used since cash-rich or more profitable firms invest more due to a higher availability of internal resources (e.g., Dobbins and Jacob, 2016; Faulkender and Petersen, 2012; Fazzari et al., 1988; Lamont, 1997). Likewise, smaller firms are expected to have better opportunities for investment (e.g., Carpenter and Petersen, 2002; Dobbins and Jacob, 2016). To measure growth opportunities, I also include Ln(Sales Growth) and Tobin's q. (e.g., Dobbins and Jacob, 2016; Jacob et al., 2018). Besides, a dummy for

effects. Firm fixed effects α_i , for instance, capture time-invariant factors at the firm level which potentially affect investment behaviour (e.g., Dobbins and Jacob, 2016; Jacob et al., 2018). Likewise, I include [group]-industry-year fixed effects $\alpha_{g,k,t}$, where [group] is a substitute for the deficit-interest-payment-cluster and individual industries are denoted by the subscript k ²⁸. Hence, firms experiencing a personal tax change in country j are compared to a control group which is operating in the same industry k and subject to similar economic conditions in terms of budget deficit and interest payments in year t . Since firms in country j are subject to the same tax system, my baseline regression employs heteroskedasticity-robust standard errors clustered at the country level.

Recalling hypothesis one in section 2, I expect the aggregate effect of personal taxes on capital investment to be ambiguous. That is, although an increase in personal taxes unambiguously increases the factor price of labour, thus making labour relatively more unattractive, capital investment of firms can respond in two ways. First, firms could treat labour and capital as complements. Thus, firms would reduce capital investment analogously to the more expensive factor labour to maintain their optimal input factor mix as determined by their production function (Dobbins and Jacob, 2016). Second, previous studies demonstrated that labour and capital can be substitutes on the margin (e.g., Dyreng et al., 2017). That is, firms partially substitute the more expensive factor labour by capital, and hence increase their capital investment even though taxes increase²⁹. I thus make no prediction on the sign of my coefficient β_1 as $\beta_1 < 0$ and $\beta_1 > 0$ are both plausible.

Table 3 presents my baseline results. In column (1), I use the top marginal income tax rate on labour income as my independent variable of interest. Columns (2) to (5) employ extended definitions of the personal tax rate which include social security contributions. Surprisingly, capital investment responses depend on the definition of personal taxes. That is, although coefficients of personal taxes are mostly positive across all five specifications, only the coefficient of the pure personal tax rate (hereafter: pure tax rate) is significant³⁰. Vice versa, all specifications including social security contributions on average have no effect on firm-level investment due to insignificant coefficients. These results have two implications. First, for the pure tax rate, my results confirm empirical findings of prior studies (e.g., Dyreng et al., 2017)

losses is added to respect that firms with negative pre-tax income are likely to invest less (Dobbins and Jacob, 2016).

²⁸My baseline results are not robust to replacing [group]-industry-year fixed effects by region-industry-year fixed effects and income-group-industry-year effects. Please refer to Table 7 in section 6 for results.

²⁹Consistent with my hypothesis development, I abstract from productivity differences between the two input factors as corresponding estimates are difficult to obtain (e.g., Dyreng et al., 2017).

³⁰Consistent with previous literature on corporate taxes (e.g., Dobbins and Jacob, 2016), dividend taxes (e.g., Alstadsæter et al., 2017) and consumption taxes (e.g., Jacob et al., 2018), coefficients on other tax rates are almost always significant and their sign is negative.

Table 3: Baseline Results

This table presents my regression results on investment behaviour from 1999 to 2013. The dependent variable is Investment. I use five different specifications of the personal tax rate. In column (1), the top marginal income tax rate on labour income is used. In Columns (2) to (5), this definition is extended and includes social security contributions for different income classes of employees in OECD countries for the 2000-2013 period. Please refer to Table A.1 in the appendix for definitions of independent variables. I further include firm fixed effects and [group]-industry-year fixed effects in all specifications, where [group] is a substitute for the Deficit-Interest-Payment-cluster. This table also reports robust standard errors clustered at the country level in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
Personal Tax	0.0367* (0.0213)				
67% Earner		0.0003 (0.0367)			
100% Earner			0.0129 (0.0243)		
133% Earner				-0.0268 (0.0308)	
167% Earner					0.0258 (0.0327)
Corporate Tax	-0.0471 (0.0355)	-0.0958*** (0.0325)	-0.0981*** (0.0326)	-0.0970*** (0.0319)	-0.0973*** (0.0309)
Consumption Tax	-0.4256*** (0.0604)	-0.5755*** (0.0835)	-0.5758*** (0.0826)	-0.5673*** (0.0819)	-0.5823*** (0.0851)
Payout Tax	-0.0094 (0.0140)	-0.0165* (0.0095)	-0.0172* (0.0095)	-0.0172* (0.0092)	-0.0182* (0.0094)
Cash Holdings	0.0200*** (0.0068)	0.0176** (0.0068)	0.0176** (0.0068)	0.0176** (0.0068)	0.0176** (0.0068)
Profit	0.0196* (0.0105)	0.0106 (0.0103)	0.0106 (0.0103)	0.0106 (0.0103)	0.0106 (0.0103)
Leverage	-0.0438*** (0.0077)	-0.0415*** (0.0074)	-0.0415*** (0.0074)	-0.0415*** (0.0073)	-0.0415*** (0.0074)
Ln(Sales Growth)	0.0032** (0.0015)	0.0033* (0.0018)	0.0033* (0.0018)	0.0033* (0.0018)	0.0033* (0.0018)
Tobin's q	0.0020* (0.0011)	0.0032** (0.0012)	0.0032** (0.0013)	0.0032** (0.0013)	0.0032** (0.0013)
Size	-0.0186*** (0.0025)	-0.0172*** (0.0024)	-0.0172*** (0.0024)	-0.0172*** (0.0024)	-0.0172*** (0.0024)
Loss	-0.0089*** (0.0014)	-0.0082*** (0.0012)	-0.0082*** (0.0012)	-0.0082*** (0.0012)	-0.0081*** (0.0012)
Accelerated Depreciation	0.0023 (0.0024)	0.0037* (0.0019)	0.0035* (0.0018)	0.0041* (0.0022)	0.0034 (0.0020)
LCB	0.0057* (0.0031)	0.0124*** (0.0027)	0.0128*** (0.0030)	0.0120*** (0.0029)	0.0124*** (0.0026)
Group Taxation	-0.0053 (0.0074)	0.0039 (0.0070)	0.0036 (0.0072)	0.0040 (0.0070)	0.0039 (0.0074)
Progressive	-0.0035 (0.0059)	-0.0060 (0.0042)	-0.0057 (0.0039)	-0.0058 (0.0042)	-0.0054 (0.0039)
GDP Growth	0.0009*** (0.0002)	0.0011** (0.0005)	0.0011** (0.0005)	0.0011** (0.0005)	0.0011** (0.0005)
Ln(GDP per Capita)	-0.0120 (0.0165)	-0.0958*** (0.0289)	-0.0978*** (0.0287)	-0.0990*** (0.0289)	-0.0940*** (0.0288)
Inflation	0.0002 (0.0002)	0.0002 (0.0005)	0.0002 (0.0004)	0.0002 (0.0004)	0.0002 (0.0005)

(Continued)

Table 3—continued

Deficit	0.0010** (0.0004)	0.0009 (0.0007)	0.0009 (0.0007)	0.0009 (0.0007)	0.0009 (0.0007)
Interest Payments	0.2739 (0.1862)	0.5349** (0.1983)	0.5299*** (0.1796)	0.5660*** (0.1966)	0.5226*** (0.1816)
Voice and Accountability	-0.0116* (0.0063)	-0.0297** (0.0144)	-0.0295** (0.0144)	-0.0285* (0.0151)	-0.0297** (0.0144)
Political Stability	0.0056 (0.0046)	0.0195*** (0.0035)	0.0188*** (0.0033)	0.0192*** (0.0034)	0.0201*** (0.0038)
Government Effectiveness	0.0125 (0.0086)	0.0234*** (0.0055)	0.0232*** (0.0058)	0.0226*** (0.0057)	0.0235*** (0.0056)
Regulatory Quality	0.0054 (0.0077)	0.0101 (0.0073)	0.0112 (0.0072)	0.0090 (0.0080)	0.0117 (0.0078)
Rule of Law	-0.0073 (0.0120)	-0.0263* (0.0134)	-0.0268* (0.0134)	-0.0248* (0.0139)	-0.0271* (0.0138)
Control of Corruption	0.0099 (0.0060)	0.0157** (0.0069)	0.0168** (0.0080)	0.0171** (0.0072)	0.0151** (0.0069)
Observations	158,760	125,582	125,582	125,582	125,582
Firm FE	Yes	Yes	Yes	Yes	Yes
(Group)-Industry-Year FE	Yes	Yes	Yes	Yes	Yes
Adjusted R-squared	0.552	0.594	0.594	0.594	0.594

showing that labour and capital are substitutes on the margin. Second, firm-level investment responses depend on the definition of the tax wedge, and thus I cannot confirm expectations about social security contributions having a similar economic effect as the pure tax rate. Yet, I would like to caution that the second implication may result from the composition of my data since data on social security contributions are only available for OECD countries. It is therefore advisable to further test this result in future studies once additional data are available.

Based on these findings, a one-percentage-point increase in the pure tax rate on average increases capital investment by 0.037 percentage points (pp) of lagged total assets³¹ which confirms that, when abstracting from productivity differences between factors, personal taxes increase the pressure on firms to substitute labour by capital on the margin. Compared to the sample average of my dependent variable Investment, this implies a relative increase of 0.51%³². For better interpretation, I convert this relative change into an implied elasticity of 0.20³³ suggesting that personal taxes are of high economic relevance for investment decisions

³¹In Table 3, I obtain a beta of 0.0367 for the average effect of personal taxes on investment. However, all tax rates in my dataset are defined between zero and one (e.g., a rate of 37% is denoted by 0.37). Thus, I multiplied the tax rate by 100 (i.e., $0.37 * 100 = 37$) to interpret the beta with respect to a one-percentage-point increase in the tax rate (e.g., from 37% to 38%). Simultaneously, I divided my beta by 100 to keep the term $\beta_1 * \text{Personal Tax}_{j,t}$ constant, thus obtaining a transformed beta of 0.000367 which equals an average change in investment by 0.0367pp.

³²The relative change of investment is defined as the quotient (β_1 / Inv_μ) of the transformed beta (i.e., 0.000367) and the sample average of Investment. In numbers, this implies $(0.000367 / 0.0719) * 100\% = 0.51\%$.

³³The implied elasticity is defined as the percentage change of investment over the percentage change of the personal tax rate ($\% \Delta \text{Inv} / \% \Delta \text{Personal}$

even though their magnitude is, in absolute terms, smaller compared to corporate taxes (between -0.4 and -0.5 as in Giroud and Rauh (2017)) and consumption taxes (between -0.24 and -0.29 as in Jacob et al. (2018); numbers are based on the draft from December 2017). Besides, since the sign of β_1 is positive, the effect of personal taxes on firm-level investment works in the opposite direction compared to other tax rates³⁴.

Overall, my baseline results confirm that on average labour and capital are substitutes on the margin even though an effect is only observed for the pure tax rate. In the following, I therefore test for cross-sectional variation in investment responses due to differences among firms in their (a) market power vis-à-vis stakeholders, (b) substitutability of input factors, and (c) financial constraints to check whether my baseline results also hold for hypotheses two, three, and four.

5.2. Cross-Sectional Variation I: Market Power vis-à-vis Stakeholders

Based on Dyreng et al. (2017) and Jacob et al. (2018), variation in capital investment responses can result from differences in the relative market power of firms, and thus their ability to shift away the economic burden of personal taxes from their shareholders. This can be explained by different labour supply (market demand) elasticities faced by firms.

(Tax). Following Jacob et al. (2018), I therefore divide the relative effect by the percentage increase of the personal tax rate. In numbers, this implies $0.51\% / (0.01 * 100\% / 0.3972)$ where 0.3972 is the sample average of the pure personal tax rate.

³⁴For completeness, relative effects and implied elasticities of all five specifications are presented in Table A.2 in the appendix.

Table 4: Market Power vis-à-vis Stakeholders, Low Profit Margin

This table presents my regression results on investment behaviour from 1999 to 2013. The dependent variable is Investment. I use five different specifications of the personal tax rate. In columns (1) and (2), the top marginal income tax rate on labour income is used. In columns (3) to (10), this definition is extended and includes social security contributions for different income classes of employees in OECD countries for the 2000-2013 period. Please refer to Table A.1 in the appendix for definitions of independent variables. Further, I interact each tax policy variable with a dummy equal to one if a firm is below the median EBIT-to-sales ratio in a given country-year (Low Profit Margin). I also include firm fixed effects and [group]-industry-year fixed effects, where [group] is a substitute for the Deficit-Interest-Payment-cluster. For all even specifications, [group]-industry-year fixed effects are replaced by country-industry-year fixed effects. This table also reports robust standard errors clustered at the country level (columns (1) and (2)) or country-industry level (columns (3) to (10)) in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Notes: In columns (3) to (10), standard errors are clustered at the country-industry level to ensure a sufficient number of clusters. Besides, I omit the variable Interest Payments in all even columns due to collinearity (i.e., columns in which the country-industry-year fixed effect is used).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Personal Tax	0.0242 (0.0226)									
Personal Tax × Low Profit Margin	0.0267** (0.0119)	0.0259** (0.0113)								
67% Earner			-0.0011 (0.0291)							
67% Earner × Low Profit Margin			-0.0090 (0.0097)	-0.0134 (0.0099)						
100% Earner					0.0262* (0.0153)					
100% Earner × Low Profit Margin					-0.0216** (0.0091)	-0.0300*** (0.0092)				
133% Earner							-0.0241 (0.0245)			
133% Earner × Low Profit Margin							0.0041 (0.0116)	-0.0013 (0.0118)		
167% Earner								0.0406* (0.0221)		
167% Earner × Low Profit Margin								-0.0200** (0.0092)	-0.0229** (0.0094)	
Observations	157,967	157,261	124,888	124,588	124,888	124,588	124,888	124,588	124,888	124,588
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
(Group)-Industry-Year FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Country-Industry-Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R-squared	0.556	0.561	0.599	0.602	0.599	0.602	0.599	0.602	0.599	0.602

That is, the more elastic a firm's labour supply (market demand), the lower the ability to pass on the personal tax incidence to workers (consumers), and hence the more of the economic burden of personal taxes is borne by firms, and ultimately, shareholders. This likely translates into higher pressure to substitute the more expensive factor labour by capital.

Hence, higher personal taxes exert higher pressure on profits, and thus are expected to affect investment responsiveness more strongly if firms have low market power. Following previous literature, I proxy a firm's market power by its EBIT margin³⁵ (e.g., Jacob et al., 2018; Lerner, 1934) and add the dummy Low Profit Margin which is equal to one if a firm is below the median EBIT-to-sales ratio in a given country-year³⁶ (e.g., Jacob et al., 2018). I subsequently interact Low Profit Margin with all tax policy variables (i.e., Personal Tax_{j,t} and control vector $T_{j,t}$) to infer whether firms with low market power respond more strongly compared to the average investment response.

Results are presented in Table 4. Interestingly, it appears that investment responsiveness of firms with low market power is ambiguous and varies with the definition of the tax wedge, too. In column (1), for instance, the interaction with the pure tax rate has a positive coefficient suggesting that firms with low market power increase their capital investment by 0.027pp of lagged total assets more strongly after an increase in personal taxes compared to the average investment response³⁷. In relative terms, this corresponds to a substantial increase in investment responsiveness by 110%³⁸ relative to the average investment response if firms have low market power. Consequently, it seems that firms facing highly elastic labour supply (market demand) bear more of the economic burden of personal taxes through lower profits, and thus are exposed to higher pressure to substitute labour by capital, which confirms my hypothesis.

Surprisingly, however, the direction of the marginal effect mostly reverses if social security contributions are included in the definition of the tax wedge, thereby contradicting my hy-

pothesis. The negative coefficients in columns (5) and (9) indicate that investment of firms with low market power responds less strongly by 0.022pp and 0.020pp of lagged total assets, respectively, compared to the average investment response. This equals a considerable decrease in investment responsiveness by 82% and 49% if firms have employees earning the average wage and 167% of the average wage, respectively. It therefore appears that, once social security contributions are considered, reduced profits translate into less resources available for investment, and thus capital investment of firms with low market power responds less strongly (Jacob et al., 2018). Yet, the negative marginal effect cannot be generalised to all income classes of employees since interaction terms in columns (3) and (7) are insignificant³⁹.

Finally, all results hold when comparing high- versus low-margin firms within the same industry in the same country (i.e., by replacing deficit-interest-payment-cluster-industry-year fixed effects by country-industry-year fixed effects). Thus, my results are likely not caused by "broader policy changes ... or other unobservable characteristics [within industry k] in ... country [j in year t]" (Jacob et al., 2018, p.21). To conclude, relative market power determines the personal tax incidence borne by firms and consequently their capital investment responses to personal tax changes. Yet, investment responsiveness of firms with low market power is ambiguous and depends on the definition of the tax wedge. That is, if the pure tax rate is used, firms with low relative market power show stronger investment responsiveness to personal tax changes compared to the average investment response. Conversely, investment of firms with low market power mostly responds less strongly compared to the average investment response once social security contributions are considered in the tax wedge⁴⁰.

5.3. Cross-Sectional Variation II: Substitutability of Labour and Capital

Although an increase in personal taxes increases the factor price of labour, and thus the pressure to substitute labour by capital, the degree of factor substitutability likely varies across firms (e.g., Dyreng et al., 2017). I therefore examine cross-sectional variation in investment responsiveness due to differences in the substitutability of input factors. There are two explanations for this phenomenon. First, the more knowledge-intensive the factor labour, the more difficult it is to substitute labour by capital since knowledge-intensive labour (e.g., R&D) is mostly difficult to automate. Thus, firms with knowledge-intensive labour are expected to substitute labour by capital to a lower degree, if at all. Second, the higher the importance of an input factor in a firm's production function, and hence for the generation of output,

³⁵I acknowledge that labour supply (market demand) elasticity is influenced by factors such as education of workers (availability of substitutes) (e.g., Dyreng et al., 2017; Jacob et al., 2018) which could serve as alternative proxies for market power. However, I abstracted from these factors for two reasons. First, these factors are not available in my dataset. Second, a firm's profit margin can be interpreted as the result of market power and thus is a conceptually correct proxy.

³⁶Alternatively, I could identify low-margin firms within the same industry using a country-industry-year distribution for Low Profit Margin. However, this would marginalise firms with relatively low profit margins in high-margin industries as low-margin firms although, in absolute terms, they are high-margin firms and vice versa. Thus, I ignore differences in the profitability of firms within the same industry and only account for differences in profitability within the same country. This also applies to subsequent tests in sections 5.3 and 5.4.

³⁷This finding is in line with previous studies suggesting that labour and capital can be substituted on the margin (e.g., Dyreng et al., 2017).

³⁸I compute the relative effect to better interpret the magnitude of the marginal effect. Following Jacob et al. (2018), the relative effect is defined as the combined effect (i.e., average plus marginal effect) over the average effect minus one. In numbers this implies $((0.0267+0.0242)/0.0242) - 1$. This calculation also applies to all other relative effects presented in subsequent analyses.

³⁹Results for all interaction terms are robust to using a tercile or quartile split. See excel file 4. Market Power Results Edited.xls for results.

⁴⁰Please note that I did not interpret average effects in this section since my research question exclusively examines whether low-margin firms respond differently from the average investment response. This also applies to subsequent analyses in sections 5.3 and 5.4.

Table 5: Substitutability of Input Factors, High Capital-to-Output Ratio

This table presents my regression results on investment behaviour from 1999 to 2013. The dependent variable is Investment. I use five different specifications of the personal tax rate. In columns (1) and (2), the top marginal income tax rate on labour income is used. In columns (3) to (10), this definition is extended and includes social security contributions for different income classes of employees in OECD countries for the 2000-2013 period. Please refer to Table A.1 in the appendix for definitions of independent variables. Further, I interact each tax policy variable with a dummy equal to one if a firm is in the top quartile of the (net-)PPE-to-sales distribution in a given country-year (High K-to-Output). I also include firm fixed effects and [group]-industry-year fixed effects, where [group] is a substitute for the Deficit-Interest-Payment-cluster. For all even specifications, [group]-industry-year fixed effects are replaced by country-industry-year fixed effects. This table also reports robust standard errors clustered at the country level (columns (1) and (2)) or country-industry level (columns (3) to (10)) in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Notes: In columns (3) to (10), standard errors are clustered at the country-industry level to ensure a sufficient number of clusters. Besides, I omit the variable Interest Payments in all even columns due to collinearity (i.e., columns in which the country-industry-year fixed effect is used).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Personal Tax	0.0493** (0.0229)									
Personal Tax × High K-to-Output	-0.0447 (0.0482)	-0.0491 (0.0492)								
67% Earner			-0.0166 (0.0287)							
67% Earner × High K-to-Output			0.0718*** (0.0206)	0.0565*** (0.0213)						
100% Earner					0.0041 (0.0152)					
100% Earner × High K-to-Output					0.0424** (0.0176)	0.0281 (0.0184)				
133% Earner							-0.0421* (0.0246)			
133% Earner × High K-to-Output							0.0598*** (0.0219)	0.0463** (0.0228)		
167% Earner									0.0111 (0.0223)	
167% Earner × High K-to-Output									0.0733*** (0.0190)	0.0591*** (0.0194)
Observations	157,999	157,296	124,897	124,598	124,897	124,598	124,897	124,598	124,897	124,598
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
(Group)-Industry-Year FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Country-Industry-Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R-squared	0.559	0.564	0.602	0.604	0.601	0.604	0.602	0.604	0.602	0.604

the more difficult it is to substitute this input factor. For instance, if labour (capital) is highly productive and therefore important for the generation of output, firms likely show a smaller (greater) substitution response from labour to capital as labour is relatively more difficult (easier) to be substituted by capital on the margin. Thus, I expect firms to show a smaller (greater) investment response if labour is knowledge intensive (capital is of high importance for the generation of output). Since proxies for labour suffer from poor coverage⁴¹, I limit the empirical analysis to my second prediction and define the dummy High K-to-Output⁴² which is equal to one if a firm is in the top quartile of the net-PPE-to-sales distribution in a given country-year. Subsequently, I interact each tax policy variable with High K-to-Output to estimate whether investment of firms responds more strongly if capital is important for the generation of output.

Table 5 presents my empirical results. Interestingly, the interaction term of High K-to-Output and the pure tax rate is insignificant whereas all specifications which include social security contributions in the tax wedge show (highly) significant, positive interaction terms. Thus, when including social security contributions in the personal tax wedge, responsiveness of capital investment increases between 0.042pp and 0.073pp of lagged total assets (depending on the income class of employees) compared to the average investment response if capital is important for the generation of output in firms, which is in line with my hypothesis. The economic magnitude of this is substantial for two reasons. First, in columns (5) and (9), investment responds more strongly to personal tax changes by factor ten and almost factor seven, respectively, if capital is important for output generation. Second, the positive marginal effect outweighs the negative average effect in columns (3) and (7) which results in a positive net effect of 0.055pp and 0.018pp of lagged total assets, respectively⁴³. Finally, the significance and magnitude of my results are mostly robust if country-industry-year fixed effects are included, and thus unobservable country-industry-year characteristics likely do not influence my results (Jacob et al., 2018). The sole difference in this case is that the marginal effect in column (6) is about half the magnitude and thus insignificant⁴⁴.

To summarise, the positive marginal effect is consistent

⁴¹My baseline sample only has 128,016 observations for R&D expenditure if values for the personal tax rate are not missing. Using R&D expenditure as a proxy for labour could therefore limit the interpretation of results towards a smaller subset of firms.

⁴²High K-to-Output is an alternative proxy for the importance of capital for output generation since estimates for factor productivities are difficult to obtain. Intuitively, I assume that a high proportion of fixed assets on a firm's balance sheet corresponds to a greater importance of capital in the production function, and thus for the generation of output. Yet, factor decisions are based on the ratio of factor productivity to factor price, and hence I acknowledge that it is conceptually reasonable but practically difficult to include a proxy for factor productivity.

⁴³I did not compute the relative effect in this case since the marginal effect outweighs the average effect, and thus the relative effect cannot be interpreted. Instead, I present the net effect which equals the sum of the average effect and the marginal effect.

⁴⁴Results are robust to using a tercile split but differ if a median split is

with my hypothesis when including social security contributions in the tax wedge, and the magnitude of this marginal effect is substantial. In other words, responsiveness of capital investment increases if capital is of high importance for the generation of output in firms. Yet, my hypothesis does not hold for the pure tax rate since, for this definition of the personal tax wedge, investment does not respond differently compared to the average investment response if capital is important for the generation of output in firms.

5.4. Cross-Sectional Variation III: Financial Constraints

Recalling hypothesis four in section 2, cross-sectional variation in investment responses can arise from "differences in the availability of internal funds" (Jacob et al., 2018, p.5). That is, if internal cash flows are the marginal source of finance for investments, investments in cash-constrained firms are likely more prone to decreases in internal cash flows than investments in cash-rich firms (e.g., Dobbins and Jacob, 2016; Faulkender and Petersen, 2012; Fazzari et al., 1988; Jacob et al., 2018). In other words, when personal taxes increase, internal cash flows, and thus the availability of resources for investment is expected to decrease more strongly if firms face financial constraints and heavily use internal funds for investments (Jacob et al., 2018). Hence, despite higher pressure to substitute labour by capital, these firms are expected to show a more negative investment response (i.e., lower investment levels) compared to the average investment response. Based on Jacob et al. (2018), I include the dummy Low Cash Flow in my regression which is equal to one if a firm is in the bottom quartile of the cash-holdings-to-total-assets distribution in a given country-year. I also interact each tax policy variable with Low Cash Flow as done in previous analyses.

Results in Table 6 indicate that investment responsiveness of financially constrained firms depends on the definition of the tax wedge as well. For example, the interaction term in column (1) has a positive but insignificant coefficient implying that investment of financially constrained firms does not respond differently than the average investment response if the tax wedge only comprises the pure tax rate, and thus my hypothesis does not hold for this specification⁴⁵. Contrarily, when including social security contributions in the tax wedge, interaction term coefficients are mostly negative and, in columns (3) and (7), significant. Thus, if firms employ workers earning 67% and 133% of the average wage, investment of financially constrained firms responds more negatively by 0.026pp and 0.024pp of lagged total assets, respectively, compared to the average investment response, which

used. In the latter case, interaction term coefficients are mostly positive but only significant in columns (9) and (10). This indicates that the median may not be an ideal threshold value. See excel file 5. Substitutability Results Edited.xls for results.

⁴⁵Interestingly, the average effect coefficient of the pure tax rate in Table 6 has a similar magnitude as in my baseline specification, and this result holds for all quantile splits. Besides, the p-value of this average effect is 0.107, and thus close to being significant at the 10% level.

Table 6: Financial Constraints, Low Cash Flow

This table presents my regression results on investment behaviour from 1999 to 2013. The dependent variable is Investment. I use five different specifications of the personal tax rate. In columns (1) and (2), the top marginal income tax rate on labour income is used. In columns (3) to (10), this definition is extended and includes social security contributions for different income classes of employees in OECD countries for the 2000-2013 period. Please refer to Table A.1 in the appendix for definitions of independent variables. Further, I interact each tax policy variable with a dummy equal to one if a firm is in the bottom quartile of the cash-holdings-to-total-assets distribution in a given country-year (Low Cash Flow). I also include firm fixed effects and [group]-industry-year fixed effects, where [group] is a substitute for the Deficit-Interest-Payment-cluster. For all even specifications, [group]-industry-year fixed effects are replaced by country-industry-year fixed effects. This table also reports robust standard errors clustered at the country level (columns (1) and (2)) or country-industry level (columns (3) to (10)) in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Notes: In columns (3) to (10), standard errors are clustered at the country-industry level to ensure a sufficient number of clusters. Besides, I omit the variable Interest Payments in all even columns due to collinearity (i.e., columns in which the country-industry-year fixed effect is used).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Personal Tax	0.0345 (0.0211)									
Personal Tax × Low Cash Flow	0.0097 (0.0155)	0.0086 (0.0157)								
67% Earner			0.0049 (0.0296)							
67% Earner × Low Cash Flow			-0.0255** (0.0114)	-0.0293** (0.0121)						
100% Earner					0.0115 (0.0158)					
100% Earner × Low Cash Flow					0.0053 (0.0113)	0.0043 (0.0121)				
133% Earner							-0.0218 (0.0247)			
133% Earner × Low Cash Flow							-0.0244* (0.0134)	-0.0273* (0.0142)		
167% Earner								0.0257 (0.0222)		
167% Earner × Low Cash Flow								-0.0049 (0.0103)	-0.0051 (0.0108)	
Observations	158,748	158,044	125,577	125,278	125,577	125,278	125,577	125,278	125,577	125,278
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
(Group)-Industry-Year FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Country-Industry-Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R-squared	0.552	0.558	0.594	0.597	0.594	0.597	0.594	0.597	0.594	0.597

Table 7: Robustness to Different Fixed Effects and Inclusion of Additional Controls

This table presents the robustness of my regression results on investment behaviour from 1999 to 2013. The dependent variable is Investment. I use five different specifications of the personal tax rate. In columns (1) and (2), the top marginal income tax rate on labour income is used. In columns (3) to (10), this definition is extended and includes social security contributions for different income classes of employees in OECD countries for the 2000-2013 period. Please refer to Table A.1 in the appendix for definitions of independent variables. I include firm fixed effects in all specifications but replace [group]-industry-year fixed effects by region-industry-year fixed effects (even columns) and income-group-industry-year fixed effects (odd columns). This table also reports robust standard errors clustered at the country level in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Personal Tax	0.0090 (0.0286)	-0.0025 (0.0202)								
67% Earner			0.0057 (0.0316)	-0.0006 (0.0326)						
100% Earner					0.0081 (0.0194)	-0.0055 (0.0143)				
133% Earner							0.0175 (0.0288)	-0.0428 (0.0261)		
167% Earner									0.0753** (0.0355)	-0.0290 (0.0277)
Observations	158,831	158,834	125,650	125,671	125,650	125,671	125,650	125,671	125,650	125,671
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Additional Country Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
(Group)-Industry-Year FE	No	No	No	No	No	No	No	No	No	No
Region-Industry-Year FE	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Income-Group-Industry-Year	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R-squared	0.552	0.550	0.592	0.590	0.592	0.590	0.592	0.590	0.592	0.590

confirms my hypothesis. The economic magnitude of this is substantial in two ways. First, investment responsiveness of financially constrained firms increases by 112% relative to the negative average effect in column (7). Second, in column (3), the negative marginal effect considerably outweighs the positive average effect yielding a combined negative net effect of -0.021 pp of lagged total assets. Like in section 5.3, however, these results cannot be generalised across all income classes of employees since interaction terms in columns (5) and (9) are insignificant, and thus investment of financially constrained firms does not respond differently than the average investment response in these specifications⁴⁶. Lastly, the significance and magnitude of my results are robust if country-industry-year fixed effects are included, and thus my results are not caused by unobservable country-industry-year characteristics (Jacob et al., 2018).

To conclude, financially constrained firms only respond more negatively if social security contributions are considered in the tax wedge, but this result cannot be generalised to all income classes of employees. Thus, my results are partially in line with results obtained for corporate taxes (e.g., Dobbins and Jacob, 2016) and consumption taxes (e.g., Jacob et al., 2018).

6. Robustness of Baseline Results

There are two main concerns about my baseline results. First, it could be argued that these are driven by the choice of my control group (i.e., comparable countries) in my cross-country panel. In other words, the deficit-interest-payment-cluster derived in my linear probability model may not be an ideal fixed effect although it compares firms in one country to a control group of firms in countries with similar economic conditions in terms of budget deficit and interest payments. Second, one could doubt whether my baseline specification accounts for all relevant variables which have an impact on firm-level investment. Thus, my baseline results could arguably suffer from omitted variable bias. To address these concerns, I modify my baseline regression in two ways. First, I define control groups differently by replacing deficit-interest-payment-cluster-industry-year fixed effects by region-industry-year fixed effects and income-group-industry-year fixed effects. Doing so, I group comparable countries in seven geographic regions and four income groups as defined by the World Bank. Second, I include Openness as an additional country-level control which I previously omitted in my baseline due to poor coverage.

Table 7 displays results of my robustness test. Overall, it appears that my baseline results are not robust to modifications of fixed effects, and thus, highly dependent on the set of comparable countries used as a control group. Specifically, I observe two patterns. First, when employing region-industry-year fixed effects, all coefficients are positive, but

their magnitude and significance differ substantially from my baseline results. For instance, magnitudes of the average effect of the pure tax rate and the 100% Earner have substantially decreased, and the average effect of the pure tax rate becomes insignificant. Contrarily, all other definitions of personal taxes become (substantially) more positive, and in case of the 167% Earner, statistically significant. Thus, it can still be proved that labour and capital are substitutes on the margin, but this only holds statistically for the 167% Earner. Second, the sign of β_1 fully reverses (i.e., turns negative) for all personal tax definitions when using income-group-industry-year fixed effects. However, the hypothesis that labour and capital are complements on the margin cannot be proved since all coefficients are statistically insignificant across all specifications in which income-group-industry-year fixed effects are used.

Importantly, these results are not driven by the inclusion of Openness for two reasons. First, results are similar in magnitude and significance if additional country-level controls are not included but fixed effects are substituted by region- and income-group-industry-year fixed effects. Second, even if deficit-interest-payment-cluster-industry-year fixed effects are not replaced by alternative fixed effects, baseline results are robust to the inclusion of Openness⁴⁷. To conclude, my baseline results appear to be ambiguous since I cannot eliminate concerns that these are potentially driven by the definition of my control group.

7. Conclusion

In this thesis, I present empirical evidence on the effect of personal taxes on firm-level investment by exploiting personal tax changes in my international panel data of 115 countries from 1999 to 2013. My findings are based on a linear regression model in which five different definitions of the personal tax wedge are regressed against capital investment of firms. Interestingly, my results show that investment responses differ depending on the definition of the personal tax wedge. In my baseline regression, firms on average show a positive capital investment response if personal taxes increase, but this effect can only be validated for the pure personal tax rate. Likewise, I obtain mixed results when testing for cross-sectional variation in capital investment responses due to differences in relative market power, the ability to substitute input factors, and financial constraints. My baseline results also vary strongly depending on the control group used, and thus are not robust to the inclusion of different fixed effects.

The positive average capital investment response can be explained by the higher substitution pressure faced by firms. That is, if firms bear part of the economic burden of personal taxes, an increase in personal taxes, *ceteris paribus*, increases

⁴⁶Results are robust to using a tercile or median split. Please refer to the excel file 6. Fin Constraints Results Edited.xls for detailed results.

⁴⁷Tables are not included in this thesis. Please refer to the excel files 3. Baseline Results Edited.xls and 7. Robustness FE Controls Edited.xls for detailed results.

the factor price of labour, and thus exerts higher pressure on corporate profits. Profit-maximising firms therefore counteract this pressure by (partially) substituting the more expensive input factor labour by capital. This mechanism, however, does not explain why including social security contributions in the personal tax wedge triggers a different capital investment response of firms compared to the pure personal tax rate. Yet, since the composition of my data may have caused this difference, it is advisable to test this result in future studies once additional data on social security contributions are available.

Eventually, my results have one potential implication for managers and policy makers. That is, personal taxes increase the factor price of labour, and thus affect decisions on the optimal input factor mix of firms. Further, when abstracting from productivity differences between factors, this 'price increase' likely discriminates the input factor labour, and thus labour-intensive firms, while favouring the factor capital, and thus capital-intensive firms. However, since input factor decisions are also a function of input factor productivities, my results cannot fully confirm this prediction since they only consider changes in the factor price. It is therefore reasonable to include estimates for labour and capital productivity in future studies before making clear policy recommendations and reform proposals. Thus, the results of this thesis can rather be understood as a first step towards reaching this goal and need to be further investigated in future theses.

References

- Alesina, A., Ardagna, S., Perotti, R., and Schiantarelli, F. Fiscal policy, profits, and investment. *American economic review*, 92(3):571–589, 2002.
- Alstadsæter, A. and Jacob, M. *Income shifting in Sweden: an empirical evaluation of the 3: 12 rules*. Ministry of Finance, Regeringskansliet, 2012.
- Alstadsæter, A., Jacob, M., and Michaely, R. Do dividend taxes affect corporate investment? *Journal of Public Economics*, 151:74–83, 2017.
- Arnold, J. M., Brys, B., Heady, C., Johansson, Å., Schweltnus, C., and Vartia, L. Tax policy for economic recovery and growth. *The Economic Journal*, 121(550):F59–F80, 2011.
- Asker, J., Farre-Mensa, J., and Ljungqvist, A. Comparing the investment behavior of public and private firms. Technical report, National Bureau of Economic Research, 2011.
- Auerbach, A. J., Aaron, H. J., and Hall, R. E. Corporate taxation in the united states. *Brookings Papers on Economic Activity*, 1983(2):451–513, 1983.
- Baker, M., Stein, J. C., and Wurgler, J. When does the market matter? stock prices and the investment of equity-dependent firms. *The Quarterly Journal of Economics*, 118(3):969–1005, 2003.
- Becker, B., Jacob, M., and Jacob, M. Payout taxes and the allocation of investment. *Journal of Financial Economics*, 107(1):1–24, 2013.
- Carpenter, R. E. and Petersen, B. C. Is the growth of small firms constrained by internal finance? *Review of Economics and statistics*, 84(2):298–309, 2002.
- Chetty, R. and Saez, E. Dividend and corporate taxation in an agency model of the firm. *American Economic Journal: Economic Policy*, 2(3):1–31, 2010.
- Cummins, J. G., Hassett, K. A., and Hubbard, R. G. Tax reforms and investment: A cross-country comparison. *Journal of public Economics*, 62(1-2): 237–273, 1996.
- Djankov, S., Ganser, T., McLiesh, C., Ramalho, R., and Shleifer, A. The effect of corporate taxes on investment and entrepreneurship. *American Economic Journal: Macroeconomics*, 2(3):31–64, 2010.
- Dobbins, L. and Jacob, M. Do corporate tax cuts increase investments? *Accounting and Business Research*, 46(7):731–759, 2016.
- Dreßler, D. and Overesch, M. Investment impact of tax loss treatment—empirical insights from a panel of multinationals. *International Tax and Public Finance*, 20(3):513–543, 2013.
- Dyregang, S. D., Jacob, M., Jiang, X., and Müller, M. A. Tax avoidance and tax incidence. *Unpublished Working Paper*, 2017.
- Faulkender, M. and Petersen, M. Investment and capital constraints: repatriations under the american jobs creation act. *The Review of Financial Studies*, 25(11):3351–3388, 2012.
- Fazzari, S., Hubbard, R. G., and Petersen, B. C. Financing constraints and corporate investment. *Brookings Papers on Economic Activity*, 1, 141–206, 1988.
- Fuest, C., Peichl, A., and Siegloch, S. Do higher corporate taxes reduce wages? micro evidence from germany. *American Economic Review*, 108(2):393–418, 2018.
- Giroud, X. and Rauh, J. State taxation and the reallocation of business activity: Evidence from establishment-level data. *Working Paper*, 2017.
- Jacob, M., Michaely, R., and Müller, M. A. Consumption taxes and corporate investment. *Forthcoming in: Review of Financial Studies*, 2018.
- Katz, H. C. The decentralization of collective bargaining: a literature review and comparative analysis. *ILR Review*, 47(1):3–22, 1993.
- Lamont, O. Cash flow and investment: Evidence from internal capital markets. *The Journal of Finance*, 52(1):83–109, 1997.
- Lee, Y. and Gordon, R. H. Tax structure and economic growth. *Journal of public economics*, 89(5-6):1027–1043, 2005.
- Lerner, A. P. The concept of monopoly and the measurement of monopoly. *The Review of Economic Studies*, 1(3):157–175, 1934.
- Ljungqvist, A. and Smolyansky, M. To cut or not to cut? on the impact of corporate taxes on employment and income. *Finance and Economics Discussion Series*, pages 1–60, 2016.
- Organisation for Economic Co-operation and Development. Revenue Statistics 2017: Tax revenue trends in the OECD. Retrieved on May 17, 2017. URL <https://www.oecd.org/tax/tax-policy/revenue-statistics-highlights-brochure.pdf>.
- Pindyck, R. S. and Rubinfeld, D. L. *Microeconomics* (9. ed., Global ed., The Pearson series in economics). Harlow: Pearson, 2018.
- Yagan, D. Capital tax reform and the real economy: The effects of the 2003 dividend tax cut. *American Economic Review*, 105(12):3531–3563, 2015.



Living is easy with eyes closed - Strategische Unwissenheit und eigennütziges Verhalten

Christiane Czech

Karlsruher Institut für Technologie

Abstract

The issue of avoiding information about the consequences of one's own actions is discussed intensively. Acting that way, makes it harder to be judged for one's decisions. My bachelor thesis deals with strategic ignorance and self-serving behaviour. This paper aims to explore if people really avoid information to a high degree and whether there are certain situations or circumstances which influence these behaviour patterns. Four different experimental studies were used and compared to a large amount of literature. It is found that intransparency in situations allows for a moral "wiggle room" which makes people's actions more egoistic. Also, people like to be seen as altruistic. By analyzing the Bayesian signaling model which introduces an agent caring about his self-image, his economic advantages and who has the opportunity to find out about social benefits and the cost of acting social, the findings show that willful ignorance can be an excuse for selfish behaviour and helps maintain the idea that they act up to their ideals. Looking at situations where people have to bring an effort, ignorance shows better outcomes because people work harder when they don't know about the negative consequences.

Keywords: strategic ignorance; moral wiggle room; dictator games; self-serving behavior

1. Einleitung

„We don't really want to know“ (Norgaard, 2006, S. 347). So oder so ähnlich (vgl. Dawson et al., 2006, S.751) titeln viele Artikel. Sie rücken damit das Thema „Strategische Unwissenheit“ in den Fokus. Unsere Gesellschaft steht in der Kritik. Stichworte, wie der Klimawandel (vgl. Geiger und Swim, 2016, S. 88; Larsen et al., 2013, S. 148), Ausbeuterbetriebe (vgl. Paharia et al., 2013, S. 82) und Anschuldigungen von Kriegsverbrechen (Dutton et al., 2005, S. 442; Grossman, 2014, S. 2659; Sahlane, 2012, S. 459) umgeben uns wie eine Wolke. Doch was bietet uns diese kurze Phrase, die sich in unserem täglichen Sprachgebrauch mehr und mehr etabliert hat? Sie gibt uns Schutz vor dem „Was wäre wenn?“, der Möglichkeit des Wissens. Konsequenzen unseres Handelns sind nicht ersichtlich, wodurch Kritik und Bestrafung weniger Anknüpfungspunkte geboten werden. Nicht nur das Urteilsvermögen der anderen Personen wird dadurch eingeschränkt, sondern auch das Aufkommen selbstkritischer Gedanken. Wissen bringt uns emotional und ethisch in die Verantwortung (vgl. Ehrlich und Irwin, 2005, S. 266). Das macht Ignoranz besonders attraktiv. Wie schon Stuart Firestein in seinem Buch feststellte: „Ignorance follows knowledge, not the other way around“ (Firestein, 2012, S. 11). In-

formationen werden strategisch aus dem Wissen heraus vermieden, dass diese nachfolgende Handlungsentscheidungen beeinflussen könnten. Diese Aussage präzisiert den Titel meiner Arbeit. Trotz der ökonomischen und ethischen Vorteile, die Ignoranz bietet, begibt sich der aktiv unwissend Bleibende in eine moralische Zwickmühle, weshalb sich die Frage stellt, ob ein Großteil der Bevölkerung wirklich dieser Attraktivität unterliegt. „Vermeiden Individuen gezielt Informationen, die für sie unbequem sein könnten?“ Mit dem positiven Beweis dieser Forschungsfrage stellt sich einem die Überlegung: „Welcher Situationszustand fördert diese Verhaltensweise, wenn sie auftritt?“ Die von diesen Fragestellungen aufgeworfene Thematik soll in meiner Arbeit systematisch beleuchtet werden. Zunächst soll sich mit der ersten der beiden Fragen befassen werden, wozu die Forschungsergebnisse von Dana, Weber und Kuang zur näheren Analyse herangezogen werden. Ihre Arbeit befasst sich mit der Frage, ob Großzügigkeit wirklich als Beweis für das Ausleben von Fairnessidealen gesehen werden kann (vgl. Dana et al., 2007, S. 77). Mit dem Negieren dieser Angelegenheit, können sie parallel das Vermeiden von Informationen in intransparenten Situationen bei ihren Probanden nachweisen. Auf ihren Ergebnissen aufbauend, folgt die Betrachtung der Veröffentlichung

Grossmans 2014, der die von DWK gezeigte Ignoranz von Informationen bestätigen kann und weiter auf die Situationsgestaltung, bei der die Wahl von Ignoranz besonders attraktiv ist, eingeht (vgl. S. 2659). Dass nicht nur externe Faktoren bei der Wahl von Ignoranz und eigennützigem Verhalten eine Rolle spielen, können Grossman und van der Weele in ihrer gemeinsamen Arbeit 2017 feststellen, in dem sie dem Wechselspiel von finanziellen Vorteilen und dem Grad der Wichtigkeit der Selbstwirkung der einzelnen Person einen besonderen Stellenwert zuordnen und den unterschiedlichen Handlungsmotiven, die durch ihr Modell aufgeworfen werden, nachgehen (vgl. S. 173f.). Im letzten Abschnitt der Arbeit soll der real-effort Experimentaufbau unternehmensnahe Einblicke beschreiben. Wie verhalten sich Experimentteilnehmer, wenn Wissen ihr Leistungsverhalten real beeinflusst, in diesem Fall über eine als negativ empfundene Wohlfahrtsorganisation. Das verblüffende Ergebnis, dass Ignoranz auch prosozial Einfluss nehmen kann, kann die bisher eher im negativen Licht stehende Thematik mit einer neuen positiven Note abrunden.

2. Überblick über strategische Unwissenheit und eigennütziges Verhalten

Wenn in der Literatur das Thema „Unwissenheit im Kontext mit eigennützigem Verhalten“ angesprochen wird, so werden dafür Anknüpfungspunkte in unterschiedlichsten Lebensbereichen gewählt. Manche Themengebiete sind dafür prädestiniert, da sie besonders massive individuelle oder gesellschaftliche Folgen aufweisen. Ein Beispiel für strategische Ignoranz, deren Konsequenzen bei einem persönlich liegen, ist die Tatsache, dass 58% der Personen es vermeiden, den Kaloriengehalt von Lebensmitteln zu erfahren. Die Information beeinträchtigt nicht nur das aktuelle Genusserebnis, sondern verbirgt auch spätere negative Gesundheitsfolgen (vgl. Thunström et al., 2016, S. 117). Ebenso verhält es sich in Bezug auf Spenden. Hier sind die Wohltäter oftmals nicht daran interessiert den Wirkungsgrad ihrer Spende zu erfahren (vgl. Niehaus, 2013, S.3) oder vermeiden Informationen, um gar nicht erst geben zu müssen (Exley, 2015, S.612). Neueste Forschungsergebnisse zeigen jedoch, dass die Konfrontation mit unvermeidbaren Informationen die Teilnehmer nicht unbedingt zum Spenden anregen, sondern diese im Gegenzug beiläufige positive Informationen stärker gewichten lassen und als Entschuldigung verwenden um nicht spenden zu müssen (vgl. Exley und Kessler, 2017, S.25). Beim Thema „Klimaschutz“ zeigt sich, dass die gezielte Konfrontation von Probanden mit Informationen, die ebenfalls Zahlen über das Engagement anderer Personen enthalten, die Ignoranzrate von 53% auf 29% senkt (vgl. Thunström et al., 2014, S.210). Ein Teil der Bevölkerung leugnet sogar die Existenz des Klimawandels, was eine positive Einflussnahme erschwert (vgl. Hobson und Niemeyer, 2013, S. 396; vgl. Golman et al., 2017, S. 128). Damit ist der Klimawandel ein Musterbeispiel in der Fachliteratur, die Ignoranz aufzeigt und die gesellschaftlichen Folgen von Ignoranz anmahnt. Zu diesen negativen Vorkommnissen

dürfen politische und industrielle Korruption sowie Konfliktsituationen, die vom Krieg bis hin zum Völkermord reichen, eingereiht werden (vgl. Grossman, 2014, S. 2659). Einzelne Gruppen, wie Medien, Parteien und einflussreiche Unternehmen, tragen dabei aktiv zum Fortbestehen eines verzerrten und ungerechten Systems bei (vgl. Admati, 2017, S. 25). Bei näherer Betrachtung von strategischer Ignoranz ist auffällig, dass besonders in Situationen, die intransparent sind, Informationen aktiv vermieden werden (vgl. Dana et al., 2007, S. 78f.). Dabei ist zusätzlich die Art und Weise der Aufstellung von Handlungsalternativen entscheidend. Ist die Möglichkeit gegeben, sich passiv in einer undurchsichtigen Situation gegen Informationen zu entscheiden, tun dies signifikant mehr (Standardnormalverteilung: $p = 0,001$), als wenn sich der Handelnde aktiv dafür entscheiden muss (vgl. Grossman, 2014, S. 2662f.). Grossman und van der Weele können in der gemeinsamen Arbeit sogar den situationsbezogenen verhaltenspsychologischen Aspekt aufdecken: Personen wägen aufgrund der Wertschätzung ihres Selbstbilds und ihrer materiellen Vorteile ihre Informations- und Handlungsentscheidungen ab (Grossman und Van der Weele, 2017, S.206f.). Sowohl die Arbeiten von Dana, Weber und Kuang als auch von Grossman und van der Weele sind sehr abstrakt gehalten und zeigen allgemein strategische Verhaltensmuster auf, wodurch sie auf beliebige Bereiche übertragbar sind und deswegen von anderer Literatur als Grundlagenforschung verwendet werden. Insbesondere dem Feld der Wirtschaft wird Ignoranz und eigennütziges Verhalten nach gesagt. Gut bezahlte Manager geraten dabei immer wieder in den Fokus, nicht zuletzt bei der Abgasaffäre der Automobilindustrie (vgl. Claassen, 2017, S.3). Sie werden beschuldigt, Ignoranz gezielt als Leistungsansporn gegenüber ihren Mitarbeitern einzusetzen, in dem sie ihren Untergebenen eine faktische Autorität gewähren und moralische Zwiespalte reduzieren. Der Kritik an diesen Verhaltensweisen wird sich entzogen. Allgemein vermeiden es Menschen eine Rückmeldung zu ihren Verhaltensweisen zu erhalten. Dies kann durch eine stetig wachsende, auf theoretischen und empirischen Beweisen aufbauende Literatur nachgewiesen werden (vgl. Grossman, 2013, S.6; Aghion und Tirole, 1997, S.8; Crémer, 1995, S. 293f.). Die Stärke von Ignoranz und die damit verbundenen egoistischen Verhaltensweisen werden besonders verstärkt/beeinflusst: durch Abgrenzungen von Gruppen innerhalb eines Unternehmens, beim Auftreten einer Prinzipal-Agenten-Konstellation, durch die Unternehmensethik, durch Machtbefugnisse, den Erfolgsdruck und die finanziellen Anreize bzw. sektoralen Arbeitstätigkeiten (vgl. Kajackaite und Gneezy, 2017, S.433; Conrads et al., 2013, S. 6f.; Barfort et al., 2015, S. 28; Galperin et al., 2011, S. 407f.; Cartwright und Menezes, 2014, S. 58; Hamman et al., 2010, S. 1843; Pierce und Snyder, 2008, S. 1891). Diese Aspekte zeigen, dass die Grundlagenforschung zurecht moralische Anliegen des Agenten gewichtet und versucht, diese in Form von komplexen Modellen zu erklären, um Informations- und Handlungsentscheidungen sinnvoll begründen zu können (vgl. Grossman und Van der Weele, 2017, S. 177f.).

3. Erforschung des moralischen Spielraums

Wie bereits im vorherigen Kapitel erwähnt, vermeiden Menschen im Allgemeinen häufig die Informationsaufnahme, um sich nicht mit den Folgen ihrer Handlungen auseinander setzen zu müssen. Die Überlegungen, die dem zuvor Beschriebenen zugrunde liegen, sind Bestandteil der Experimentreihe von Dana, Weber und Kuang. In dem Paper „Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness“ werden die Ergebnisse der Autoren protokolliert. Beginnend, die moralischen Vorstellungen und Fairnessideale ihrer Teilnehmer zu untersuchen, stellen sie sich die Frage, ob Großzügigkeit in einem Experiment wirklich als Beweis dafür gesehen werden kann, dass Menschen sich um soziale Gerechtigkeit kümmern. Sie finden heraus, dass dreiviertel ihrer Probanden angeben, sich im Allgemeinen für eine faire Auszahlung in einer Entscheidungssituation zu interessieren. Aufbauend auf dieser Tatsache, zeigt ihre Experimentreihe, dass diese bekundeten Fairnessideale ebenfalls konsistent sind mit der Tatsache, dass sich Teilnehmer dazu gezwungen fühlen gerecht zu agieren, wenn ihre Handlung für andere öffentlich nachvollziehbar ist. Das heißt im Umkehrschluss, wenn die Möglichkeit dazu besteht Entscheidungen zu verschleiern, entscheidet sich ein signifikanter Anteil der Probanden bewusst für Informationsvermeidung hinsichtlich der Auszahlung anderer und/oder für die persönlich lukrativste Auszahlung. Dies ist im Experiment durch die Möglichkeit gegeben, Auszahlungsinformationen zu vermeiden beziehungsweise der Einbeziehung einer anderen ergebnisabhängigen unbekanntem Konstante in die Ergebniskette, von deren Entscheidung die Auszahlung final beeinflusst werden kann (vgl. Dana et al., 2007, S. 77f.).

3.1. Gezielte Vermeidung von Informationen und bewusste Unsicherheit über die Konsequenz des eigenen Handelns

Das Bedürfnis auf der einen Seite, von anderen wohlwollend und hilfsbereit wahr genommen zu werden, aber auf der anderen Seite seinen eigenen Profit maximieren zu wollen, steht zunächst offensichtlich in Kontrast zueinander. DWK schafft durch die Wahloption, in ihren Experimenten „unwissend zu bleiben“, die realistische Situation der Informationsvermeidung, um für Außenstehende sowie für sich selbst weiterhin ein positives Image wahren zu können. Vom scheinbaren Hang zu Fairness und Gerechtigkeit ist nun die Brücke zu strategischer Unwissenheit und weniger altruistischen Verhalten gespannt. DWK erarbeiten mit Hilfe ihrer Vorgehensweise die Grundlage für viele weitere Forschungsarbeiten in dem Bereich situationsbezogener Rechtfertigung. Auf meine Forschungsfrage bezogen beweisen DWK statistisch signifikant, dass Informationen, die Entscheider in einen moralischen Zwiespalt bringen und einen negativen Einfluss auf eigennütziges Verhalten haben könnten, gezielt vermieden werden (vgl. Dana et al., 2007, S. 68f.).

3.2. Vorstellung des Laborexperiments

An den Laborexperimenten von Dana, Weber und Kuang nehmen insgesamt 190 Studenten der Universität Pittsburgh

teil. Sie werden dabei nach dem Zufallsprinzip einer von vier verschiedenen Abwandlungen von modifizierten Diktator-Spielen zugeteilt. In denen müssen sie jeweils eine binäre Entscheidung zwischen einer gerechten und einer ungerichten Wohlfahrtsverteilung treffen. Die vier verschiedenen Experimente werden im Folgenden vorgestellt. Der Aufbau war zu Beginn bei jedem Versuch derselbe. Nachdem die Teilnehmer, die für die Entscheidungsexperimente bezahlt wurden, an Ort und Stelle erschienen waren, platzierte man sie vor Bildschirme. Über dieses Medium kommen alle Anweisungen. Sie werden zusätzlich laut vorgelesen. Nachdem jeder Proband eine Karte ausgefüllt hat, um seine Rolle (in Form eines Buchstabens) und seinen Übereinstimmungspartner (in Form einer Nummer) zugeteilt zu bekommen, wird er darüber aufgeklärt, dass ihm ein bzw. zwei Mitspieler zufällig anonym für ein einfaches Spiel zugeteilt werden. In diesem Spiel werden die Teilnehmer in Abhängigkeit von der Entscheidung eines Diktators bezahlt (in den meisten Experimenten Spieler „X“, im Multiple Dictator Experiment Spieler „X“ und „Y“). Im Anschluss wird durch ein kurzes Quiz getestet, ob alle Informationen von den Teilnehmern verstanden worden sind. Später werden sie über die realen Auszahlungen in Kenntnis gesetzt und ihrem speziellen Experiment zugewiesen. Sowohl der Empfänger als auch der Diktator werden vor jedem Experiment darüber informiert, dass kein Mitspieler jemals erfahren wird, wie der Diktator sich entschieden hat und ob er dies mit oder ohne Ausschluss von Auszahlungsinformationen tat. Des Weiteren befindet sich jeder Teilnehmer genau in einem der Experimente und weiß nichts über die individuellen Konditionen anderen Experimente, wodurch unterschiedliche Verhaltensweisen später auf direkte Effekte der Steuerung der Erhebungsmethoden zurück geführt werden können.

3.2.1. Baseline

Beim Baseline Experiment wird den 38 Probanden der Bildschirmausschnitt, wie in Abbildung 1 zu sehen, gezeigt. Die 19 Diktatoren (Spieler „X“) können sich nun zwischen der Alternative A und B entscheiden. Die Freigabe für die Button A oder B erfolgt jedoch erst nach Ablauf von 60 Sekunden. In diesem Zeitraum sollen/können sich die Teilnehmer die Auszahlungsalternativen genau anschauen. Nach der Wahl der Diktatoren werden die Empfänger ebenfalls nach ihrer hypothetischen Entscheidung gefragt, damit die Rollenverteilung weiterhin anonym bleibt. Im Anschluss werden die Spieler beim Verlassen des Raumes geheim bezahlt. Bis auf einzelne Abwandlungen sind die nachfolgenden Experimente, wenn nicht anders beschrieben, von ihrem Aufbau gleich (vgl. Dana et al., 2007, S. 71).

3.2.2. Hidden Information Experiment

In diesem Experiment können die 32 Diktatoren es vermeiden zu erfahren, wie die Auszahlung für den Empfänger, abhängig von ihrer Entscheidung, ausfällt. Die Diktatoren (Spieler „X“) wissen, dass sie für die Wahl von A \$6 und für B \$5 erhalten werden und dass ihre Empfänger je nach dem \$1 oder \$5 bekommen. Dies hängt jedoch vom Zufall ab. So

kann die Auszahlungsmatrix von konträren Interessen auftreten, diese entspricht den Auszahlungsoptionen im Baseline Experiment. Die gegenteilige Auszahlungsmatrix entspricht bei der Wahl von A \$6 (Diktator) und \$5 (Empfänger) bzw. \$5 und \$1 bei B. Bei der zweiten Variante haben die Spielpartner gleiche Interessen. Es werden jeweils zweimal beide Optionen von den Spielern durchgeführt, wobei jeder Option jeweils 16 Diktatoren zugeteilt wurden. Spieler „X“ kann sich nun zu Beginn entscheiden, ob er die Auszahlung aufdecken möchte (vgl. Abbildung 2). Den Spielern wird erklärt, dass die Fragezeichen durch die wahren Auszahlungsergebnisse ersetzt werden, wenn sie auf den „Aufdecken“-Button klicken. Während der Spieler „Y“ seine hypothetische Auswahl trifft, wird er außerdem gefragt, was er von dem konträren Interessenspiel erwartet. Wenn das Baseline Experiment die Fairnessideale der Spieler wirklich widerspiegelt, sollte der Anteil derer, die das Spiel aufdecken und sich dann für die gerechte Auszahlung entscheiden, genauso hoch sein, wie der Anteil der Personen, die im Baseline Experiment B gewählt haben. Variiert dieser Anteil, zeigt dies, dass Diktatoren unter Umständen nach einer Ausrede suchen, um sich nicht dazu genötigt zu fühlen, fair entscheiden zu müssen. Diese Ausflucht bekommen sie in Form der Möglichkeit, das Spiel verdeckt zu lassen, geliefert (vgl. Dana et al., 2007, S. 71 f.).

3.2.3. Multiple Dictator Experiment

In diesem Experiment wird von den Spielern ein weiterer Diktator hinzugefügt, das heißt, von den 30 Probanden sind 20 Diktatoren (Spieler „X“ und „Y“) und 10 passive Empfänger (Spieler „Z“). Es werden zwei Runden mit jeweils 15 Teilnehmern durchgeführt. Keiner der Diktatoren kann in diesem Experiment eine Entscheidung hinsichtlich eines fairen Versuchsausgangs unabhängig vom jeweils anderen Diktator treffen. Das bedeutet, die Auszahlung hängt von der Entscheidung der Spieler X und Y ab (vgl. Abbildung 3). Während die Probanden, denen die Rollen X und Y zugewiesen wurden, ihre Wahl treffen, wird den Z Spielern die Aufgabe gegeben sich zu überlegen, von welchem Entscheidungsergebnis der Diktatoren sie ausgehen. Das Hinzufügen eines zweiten Diktators verhindert nicht den fairen Ausgang des Experiments, da beide Diktatoren sich für A entscheiden müssen (\$6, \$6, \$1). Wenn ein Diktator sich eine gerechte Auszahlung wünscht (\$5, \$5, \$5), kann er allein mit der Wahl von B dies für alle Beteiligten entscheiden. Auch bei diesem Experiment könnte man einen gleichen Anteil an Entscheidungen für die Option B wie im Baseline Experiment erwarten. Jedoch wird durch den zweiten Diktator die Transparenz erheblich eingeschränkt, da allein die Option A zu wählen noch kein unfaires Ergebnis auslöst. Es ermöglicht es aber für einen strategischen Spieler dennoch ein egoistisches Ergebnis zu subventionieren, ohne dies direkt verantworten zu müssen (vgl. Dana et al., 2007, S. 72 f.).

3.2.4. Plausibel Deniability Experiment

Im letzten Experiment ist es den 29 Diktatoren von den 59 Teilnehmern möglich, ihre Entscheidungsgewalt abzuge-

ben. Es gibt in diesem Experiment jeweils drei Durchgänge mit je 18 Teilnehmern. Neuerungen im Vergleich zum Baseline Experiment treten in Form der Einführung eines „Cut-off“-Instruments auf. Die Teilnehmer werden darüber aufgeklärt, dass sie in einem Intervall von zehn Sekunden eine Entscheidung zu treffen haben. Diese kann ihnen allerdings von einem Zufallsgenerator, der zu einem zufälligen Zeitpunkt mit gleich hoher Wahrscheinlichkeit eine Wahl für A oder B trifft, abgenommen werden. Natürlich wird nur der Diktator später wissen, ob das „Cut-off“-Instrument Anwendung gefunden hat. Zu Beginn bekommen die Spieler eine Minute Zeit, um über ihre Auszahlungsoptionen nachzudenken, die ihnen auf dem Bildschirm (vgl. Abbildung 1) angezeigt werden. Die zehn Sekunden, in denen sie sich dann für eine Option zu entscheiden haben, starten erst mit der Bestätigung des Buttons „Spiel starten“. Auch die hypothetischen Antworten der Empfänger werden bei der Wahrscheinlichkeit, einen „Cut-off“ zu erhalten, thematisiert. Probanden, denen die Entscheidung genommen wird, werden trotzdem gefragt, wie sie sich entschieden hätten. Der „Cut-off“-Zeitpunkt wird von einer standardisierten Normalverteilung abgebildet. Ziel ist es, den Teilnehmern genügend Zeit zu geben, um ihre Wahl treffen zu können. Die Wahrscheinlichkeit, in den ersten zwei Sekunden von dem „Cut-off“-Instrument unterbrochen zu werden, liegt bei 3×10^5 . Niemand wird von der Programmierung in weniger als vier Sekunden unterbrochen. Empirische Erhebungen zeigen, dass Probanden nicht mehr als zwei Sekunden benötigten, um eine Auswahl zu treffen, nachdem sie den „Start-Button“ gedrückt haben und vorher schon eine Minute Bedenkzeit hatten. Demnach ist das „Cut-off“-Instrument eigentlich irrelevant, wenn die Probanden eine feste Präferenz gegenüber den Auszahlungen haben. Nichtsdestotrotz hilft diese besondere Methode dabei Entscheidungen zu verschleiern und dadurch zwei verschiedene Arten von moralischem Spielraum zu kreieren. Zu einen wird es dem Empfänger niemals möglich sein zu erfahren, ob die Entscheidung vom Diktator oder vom Zufallsgenerator veranlasst wurde. Im Umkehrschluss heißt das, dass den Diktator nichts davon abhält A zu wählen, außer seinem eigenen Gewissen. Zum anderen gibt es auch selbstbetrügerische Motive den moralischen Spielraum gezielt zu nutzen. Diktatoren können darauf warten, dass ihnen der Zufallsgenerator die Entscheidung abnimmt. Die Wahrscheinlichkeit, dass es zu einer gerechten Auszahlung kommt, liegt bei 50%. Tritt aber der komplementäre Fall ein, kann der Schein, nicht für das Ergebnis verantwortlich zu sein, geschickt gewahrt werden (vgl. Dana et al., 2007, S. 73 f.).

3.2.5. Auswertung scheinbarer Fairnessideale in Bezug auf Auszahlungsergebnisse

Anfangen beim Baseline Experiment entscheiden sich die Diktatoren erwartungsgemäß in der transparenten Umgebung überwiegend (74%) für die faire Auszahlungsmatrix B (\$5, \$5) (vgl. Tabelle 1). Die Empfänger entscheiden sich ohne Ausnahme für B. Diese Ergebnisse sprechen für die Großzügigkeit und Gerechtigkeitsvorstellungen der Teilnehmer, sie zeigen aber auch, wie stark sich die Teilnehmer in

einer gläsernen Situation dazu genötigt fühlen, für ein faires Ergebnis zu plädieren. Diese These kann im Hidden Information Experiment direkt statistisch signifikant belegt werden. Dort entscheiden sich in der Situation, die die gleichen Auszahlungen wie das Baseline Experiment hat (\$6, \$1), trotz der kostenlosen Möglichkeit die Auszahlungsergebnisse für den Empfänger zu erfahren, 63% für die ungerechte Option A. Der Unterschied zwischen den gewählten Optionen ist signifikant und wird mit Hilfe des Chi-Quadratstest berechnet. Dabei liegt das Ergebnis mit einem Freiheitsgrad ($\chi^2 = 4,64$) auf einem Signifikanzniveau von $p = 0,03$ (vgl. Tabelle 1). Diese Ergebnisse wurden von Larson und Noussair (2005, S. 4f.) repliziert, der in seinem Experiment eine doppelte Intransparenz einführt und die Diktatoren zum Handeln zwang. Zudem beschließen gerade einmal 56% der Teilnehmer insgesamt die Auszahlungsmatrix aufzudecken. Das zeigt, dass viele Teilnehmer konsequent ignorant gegenüber den Folgen ihres eigennützigen Verhaltens bleiben wollen. Tabelle 2 zeigt im Detail, wie sich die Probanden abhängig davon, ob sie die Spielergebnisse aufgedeckt haben, entscheiden. Wie vorher schon angedeutet, bevorzugen einige Diktatoren generell ein faires Ergebnis, woraus man schlussfolgern kann, dass Diktatoren, die die Auszahlungsmatrix aufdecken, sich ebenfalls für die faire Variante B in der Baseline Auszahlungsmatrix bzw. für A in der alternativen Auszahlungsvariante entscheiden. Doch diese Aktionskette verfolgen nur 15 von 32 Diktatoren (47%), was erneut anteilig beträchtlich weniger [$\chi^2(1) = 3,49$, $p = 0,06$] sind als im Baseline Experiment. Widergespiegelt wird dieses Ergebnis ebenfalls von den hypothetischen Entscheidungen der Empfänger. Verglichen mit dem Baseline Experiment, in dem sich noch 100% für eine faire Auszahlung entscheiden, waren es in der anderen Experimentumgebung bei einem Chi-Quadratwert von $\chi^2(1) = 3,49$ und einem Signifikanzniveau von $p = 0,001$ nur noch 59%. Dass alle (100%) Diktatoren, die im Baseline Auszahlungsfall unwissend bleiben wollten, sich für die ungerechte Auszahlung entschieden, zeigt, dass die Diktatoren im Baseline Experiment nicht unbedingt einen edelmütigen Charakter besitzen, sondern ihnen der moralische Spielraum fehlt, ihr unmoralisches Verhalten zu vertuschen (vgl. Dana et al., 2007, S. 74 f.). Das Multiple Dictator Experiment stärkt die These, dass Diktatoren, wenn sie die Chance haben, verdeckt eine ungerechte Entscheidung zu treffen, ohne für die Konsequenzen direkt verantwortlich zu sein und die Entscheidung nicht auf sie direkt zurückverfolgt werden kann, sie diese wahrnehmen. In diesem Experimentteil kommt erschwerend hinzu, dass zwar die Möglichkeit, ein gerechtes Ergebnis zu garantieren, zu jeder Zeit gewährleistet ist, aber eine egoistische Entscheidung nicht eindeutig zu einer unfairen Auszahlung führen muss. Bedingt durch diese Rahmensituation entscheiden sich 35% der Diktatoren für die faire Option B, wobei es im Baseline Fall noch 74% sind. Diese Diskrepanz ist ebenfalls statistisch signifikant mit einem Signifikanzniveau von $p = 0,02$ bei einem Chi-Quadratwert von $\chi^2(1) = 5,87$ (vgl. Tabelle 3) (vgl. Dana et al., 2007, S. 76). Die Empfänger (Spieler Z) gehen intuitiv von ähnlichen Ergebnissen aus, wie sie hier vorgefunden werden. Alle

zehn von ihnen rechnen damit, dass die Wahl eines strategischen Spielers auf die Option A fällt. In allen bisherigen Experimenten hatte geringere Transparenz Folgen zu Lasten der Empfänger, die weniger Geld erhielten. Betrachtet man nun das Plausible Deniability Experiment (vgl. Tabelle 4) sieht man, dass 7 von 29 (24%) Diktatoren in ihrer Entscheidung vom „Cut-off“-Instrument unterbrochen werden. Der durchschnittliche Eingriffszeitpunkt des Instruments liegt bei 4,30 Sekunden. Man rufe sich in Erinnerung, dass bei den Vortests keiner der Probanden länger als zwei Sekunden für seine Entscheidung benötigt hat. Das zeigt, dass ein Viertel der Teilnehmer eine Entscheidung so lange heraus zögert bis sie ihnen abgenommen wird. Davon abgesehen entscheiden sich nur noch 10 von 29 (34%) bewusst für eine faire Auszahlung. 12 der 22 (55%) nicht vom „Cut-off“-Instrument Betroffenen entscheiden sich für die egoistische Auszahlungsmatrix, was einen anteiligen Unterschied zum Baseline Experiment auf einem Signifikanzniveau von $p = 0,07$, unter Verwendung eines Chi-Quadratstests ($\chi^2(1) = 3,35$) darstellt. Wenn man sich jedoch die Diskrepanz der hypothetischen Entscheidungen der Empfänger für ein gerechtes Ergebnis (13 von 29 / 34%) im Vergleich zum Baseline Experiment (100%) anschaut, ist diese durchaus signifikant bei einem Signifikanzniveau von $p = 0,001$ ($\chi^2(1) = 15,72$). Die Empfänger hatten entsprechend den zugrunde liegenden hypothetischen Entscheidungen erneut die richtige Intuition und erfassen damit die von den Experimenten hervorgerufenen Verhaltensweisen. Schlussendlich kann DWK bilanzieren, je mehr moralischen Spielraum sie den Probanden im Versuch gewähren, desto weniger wird von den Diktatoren für ein faires Ergebnis votiert (vgl. Tabelle 5) (vgl. Dana et al., 2007, S. 77). Dabei können drei verschiedene Verhaltensmuster identifiziert werden. Zum einen Personen, die sich stets altruistisch verhalten (ca. 35%), dann Probanden, die immer egoistisch agieren (ca. 25%), und dann eine dritte Gruppe, die sich in transparenten Situationen für eine faire Auszahlung entscheidet, aber, wenn die Handlungsstränge jedoch verschleiert sind, auch eigennützig handelt. Abschließend kann DWK mit ihren Experimenten zwei entscheidende Aussagen für sich beanspruchen. Zum einen, dass, wenn ein Individuum bereit ist zu teilen, es tendenziell eher an gerechten Auszahlungen interessiert ist. Zum anderen berufen sie sich darauf, dass sie nicht beweisen wollen, dass sich Personen bewusst auf Kosten anderer selbstsüchtig bereichern, sondern, wie auch schon andere Wissenschaftler bestätigten (vgl. Murnighan/Oesch/Pillutla 1982 S. 405f.; Mitzkewitz und Nagel, 1993, S. 193; Dana et al., 2006, S. 199; Lazear et al., 2012, S. 28), unter der Grundbedingung egoistisch handeln, wenn sie keine Informationen über die Folgen ihres Handelns haben (vgl. Dana et al., 2007, S. 78 f.).

3.3. Kritische Auseinandersetzung mit den Experimentergebnissen

DWK zeigen in ihrer Experimentserie sehr anschaulich, dass sich grundsätzlich ein Drittel einer Personengruppe situationsunabhängig altruistisch verhält. Um die Probanden in eine alltagsähnliche Situation zu versetzen, verwenden

DWK mehrere Spielvarianten, im speziellen die Diktator-Spiel Variante, diese wird auch von den Wissenschaftlern des vierten und fünften Kapitels verwendet. Die Entscheidungen der Teilnehmer entstammen dadurch ihren sozialen und persönlichen Normen und Heuristiken. Statt strikt strategisch zu handeln, verhalten sich Individuen in einem Spiel eher so, wie sie sich auch in einem anderen sozialen Kontext verhalten würden. Das Diktator-Spiel spiegelt damit insbesondere Umgangsformen der Teilnehmer wider, die sie im normalen Leben erlernt haben (vgl. Murnighan und Wang, 2016, S. 89). Nach und nach beweisen sie, dass die Überlegung, dass Menschen sich dazu gezwungen fühlen zu geben, konsistent mit der These ist, dass nicht der gesamte Anteil an Probanden, die sich im Baseline Experiment für die faire Auszahlungsvariante entscheiden, diese auch wirklich wertschätzen (vgl. Dana et al., 2007, S. 77) (vgl. Tabelle 5). Besonders gut bilden diese Ergebnisse das Hidden Information Experiment und das Multiple Dictator Experiment ab. In Anbetracht dessen ist das Plausible Deniability Experiment weniger aussagekräftig. Es zeigt lediglich auf, dass sich Probanden, tendenziell davor scheuen, überhaupt eine Entscheidung zu treffen und sich somit der Verantwortung entziehen. Deswegen ist die Gegenüberstellung des Anteils an Personen, die sich für eine gerechte Auszahlungsmatrix entscheiden (34%; vgl. Tabelle 5), berechnet auf Grundlage der Gesamtdiktatoren des Experiments, auch wenig aufschlussreich, da hier ebenfalls die Teilnehmer mitzählen, die vom „Cut-off“-Instrument unterbrochen wurden und ihre Entscheidung demzufolge nicht selbst trafen. Unabhängig davon bejaht die Experimentreihe von DWK sehr passgenau die erste Forschungsthese. Ein Großteil der Probanden (37%, vgl. Tabelle 1) vermeidet gezielt Informationen, um nicht mehr geben zu müssen. Individuen suchen gezielt Situationen, die nicht transparent sind, um nicht (moralisch) gezwungen zu sein, geben zu müssen, resümiert das Paper von DWK. Diese Quintessenzen gehen nicht nur mit einander einher, sondern werden zudem noch von den Forschungsergebnissen dritter getragen (vgl. Larson und Capra, 2009, S. 467). Auch Lazear, Malmendier und Weber kommen zu dem Schluss, dass Menschen versuchen Situationen, in denen sie teilen müssten, zu vermeiden (vgl. Madrian und Shea, 2001, S. 138, Broberg et al., 2007, S. 33). Zu dieser Folgerung gelangt auch Nyborg, in deren Experiment die Probanden sogar bereit sind Geld zu zahlen, um keine Informationen zu erhalten, auf Grund derer sie vielleicht spenden müssten (vgl. Nyborg, 2011, S. 263). Diese Forschungsergebnisse finden sich in der Wirtschaft wieder, wo zum Beispiel Geschäftsführer bei anstehenden Kündigungen nicht über die persönlichen Umstände ihrer Mitarbeiter Bescheid wissen möchten. Dabei kontrolliert sie keiner, ob sie Informationen über ihre Untergebenen einholen oder aktiv unwissend über die Konsequenzen ihrer Handlung bleiben. Selbst wenn der Abbau von Arbeitsplätzen unvermeidbar ist, könnte durch Gespräche und Weiterempfehlungen an andere Firmen der Branche, die zukünftige Absicherung der Gekündigten gesichert werden und Fachwissen bewahrt werden. Dies birgt jedoch einen enormen Zeit- und Arbeitsaufwand, der gern vermieden wird. Abschließend kann festgestellt

werden, dass Individuen sich gezielt intransparente Situationen suchen, um altruistische Handlungen nicht ausführen zu müssen und im Gegenzug zum eigenen Vorteil aktiv zu handeln. Doch wie kann innerhalb einer intransparenten Situation eigennütziges Handeln positiv oder negativ beeinflusst werden? Diese Thematik wird im folgenden Kapitel genauer untersucht.

4. Strategische Ignoranz und die Stabilität sozialer Präferenzen

Probanden vermeiden es in Diktator-Spielen häufig, Informationen über die Folgen ihres Handelns zu erhalten, ganz gleich ob ihre Entscheidung für den Empfänger rentabel ist oder ihm schadet. Lieber ziehen sie es vor, egoistisch den durch ihre Ignoranz eigens geschaffenen moralischen Spielraum auszunutzen. Bisher wurde diese Verhaltensweise jedoch nur in einer Umgebung festgestellt, in der sich die Testpersonen aktiv für das Erhalten von Informationen entscheiden mussten. Im Umkehrschluss bedeutet es, dass Untätigkeit mit Ignoranz gleich zu setzen ist. Wie jedoch verhalten sich Probanden, wenn sie sich aktiv dafür entscheiden müssen ignorant zu sein oder Informationen zu erhalten oder aber ignorant zu bleiben? Dieser Fragestellung widmet sich Grossman in seiner Arbeit „Strategic Ignorance and the robustness of social preferences?“ Durch verschiedene Experimentsituationen kommen die Autoren zu dem Ergebnis, dass sich die Verhaltensweisen und die Handlungen der Diktatoren, je nach dem auf welche Art und Weise der moralische Spielraum wählbar ist, deutlich unterscheiden. Obwohl die Ausbeutung von moralischem Spielraum nicht nur ein Artefakt ist, ist er, ähnlich wie das soziale Verhalten selbst, Umwelteinflüssen und psychologischen Faktoren unterworfen, die seine Wirkung verstärken oder untergraben können. Unabhängig davon zeigen Grossmans Experimente, dass moralischer Spielraum nur dann genutzt wird, wenn er passiv wählbar ist, oder die erhaltenen Informationen erst dann eintreffen, wenn sie die Entscheidung nicht mehr tangieren können (vgl. Grossman, 2014, S. 2659f.).

4.1. Nutzung passiver Spielräume um egoistische Ergebnisse geltend zu machen

Das Paper von Grossman ist insbesondere in Bezug auf die zweite Forschungsfrage sehr passend, da es auf die Ergebnisse von DWK aufbaut und an dem Punkt „Moralischer Spielraum gegeben durch Intransparenz und selbst aufgebaute Unwissenheit“ ansetzt und diesen weiter erforscht. Die Experimentreihe von Grossman ist der Art gestaltet, dass die ersten beiden Experimente eine Replikation zweier Experimente von DWK darstellen und deren Ergebnisse ohne größere Abweichungen bestätigen. Die darauf folgenden Experimente weisen eine Steigerung der Manier auf, dass sich die Probanden, beginnend mit einer passiven Entscheidung für Unwissenheit, von Experiment zu Experiment immer aktiver dafür entscheiden müssen, keine Informationen über die Auszahlungsmatrix des Empfängers zu erhalten. Die Ergebnisse zei-

gen, dass moralischer Spielraum für sich selbst ein ganz eigenes Phänomen darstellt. Auf der einen Seite werden prosoziale Ansätze, durch die gegebene Möglichkeit in den Experimenten gegenüber den Handlungsfolgen ignorant zu bleiben, untergraben. Aber auf der anderen Seite werden prosoziale Ideale gefördert, wenn Unwissenheit aktiv gewählt werden muss (vgl. Grossman, 2014, S. 2659). Durch diesen Ansatz kann Grossman offenbaren, dass nur, wenn der moralische Spielraum passiv gewählt werden kann, dieser auch für die Vermeidung von Informationen genutzt wird.

4.2. Aufbau und Vorgehensweise der Experimente

In dem Experiment von Grossman werden Studenten der University of California zufällig als Probanden rekrutiert. Die Teilnehmer spielen dasselbe Diktator Binär Spiel, wie es in den Experimenten von DWK verwendet wurde. Sie bekommen somit auch dieselbe Einführung (vgl. Kapitel 3.2) (vgl. Grossman, 2014, S. 2660). Es gibt wieder die beiden Spielvarianten „Konträre Spielinteressen“ und „Gleiche Spielinteressen“. Nach jedem Experiment werden die Probanden gefragt, wie hoch ihrer Meinung nach der Anteil der Teilnehmer ist, von dem sie vermuten, dass sie sich wissend, dass sie sich in dem konträren Interessenspiel befinden, für die faire Variante B entscheiden bzw. nach dem Anteil, von dem sie erwarten, dass sie sich dafür entscheiden, die Auszahlungsmatrix aufzudecken. Die Antwort auf diese Fragestellung wurde, wenn sie sich in einem 5% Punkte Intervall nahe des richtigen Ergebnisses befand, mit zusätzlichen \$5 belohnt. Diese weiteren Untersuchungen sind jedoch für das nachfolgende Ergebnis nicht relevant und werden deswegen im Fortlauf auch nicht weiter aufgegriffen (vgl. Grossman, 2014, S. 2661). Wie auch bei DWK liefert das Baseline Experiment die Grundlage für die darauf folgenden Experimente. Das Baseline Experiment wird mit den Probanden genau wie bei DWK durchgeführt (vgl. Kapitel 3.2.1) (vgl. Grossman, 2014, S. 2661). Die Diktatoren spielen wissend die Spielversion mit konträren Spielinteressen, wobei die Empfänger ihre hypothetische Entscheidung treffen. Das nächste Experiment nennt sich Default NR, welches das Hidden Treatment Experiment von DWK (vgl. Kapitel 3.2.2) wiedergibt. Der einzige Unterschied ergibt sich in der Darstellung. Bei Grossman müssen die Probanden auf zwei unterschiedlichen Bildschirmbildern zunächst die Entscheidung treffen, ob sie das Spiel aufdecken wollen bzw. welche Auszahlungsoption sie wählen. Bei DWK sind beide Entscheidungen auf einem Bildschirmbild zu treffen. Den Teilnehmern bei Grossman wird bei dem ersten Entscheidungsbild die Klickoption „Spiel aufdecken“ und „Fortfahren“ zur Verfügung gestellt, wobei letzteres bereits in der Vorauswahl ist und auch mit einem dritten Knopf „Ok“ ebenfalls bestätigt werden kann. Im Active Choice Experiment spielen die Probanden dasselbe Spiel wie bei Default NR jedoch mit dem Unterschied, dass sie auf dem ersten Entscheidungsbildschirm auf die Frage „Spiel aufdecken?“ die Wahl zwischen „Ja“ und „Nein“ haben. Dabei ist keiner der beiden Buttons als Vorauswahl hinterlegt. Demzufolge ergibt sich eine neutrale Ausgangslage, in der die Probanden sich aktiv für eine Option entscheiden müssen,

um zum Auszahlungsentscheidungsbildschirm zu gelangen. Auch das Default R Experiment unterscheidet sich nur marginal von dem Default NR Experiment. Diesmal können die Probanden entweder „Spiel nicht aufdecken“ oder „Fortfahren“ anklicken, wobei „Fortfahren“ durch eine Vorauswahl bereits markiert ist und ebenfalls mit einem „Ok“-Button bestätigt werden kann. Im Strategy Method Experiment müssen sich die Probanden sowohl für eine Auszahlungsoption in der Situation von konträren als auch von gleichen Interessen entscheiden. Außerdem erfolgt bei diesem Spiel die Informationsfrage auf demselben Bildschirmbild wie die Auszahlungsentscheidung. Auf dem Monitor können die Probanden in einem separaten Fenster unterhalb des Bereichs der Auszahlungsentscheidung für Auszahlungsinformationen votieren, wobei sich unabhängig davon weder ihre Wahl noch die Auszahlung revidieren lässt. Wie auch schon bei dem Default NR Experiment kann der Vorgang auch ohne Auswahl der Informationsentscheidung fortgeführt werden. Der Button „Nicht aufdecken“ befindet sich hierbei in der Vorauswahl, der Knopf „Spiel aufdecken“ hingegen muss aktiv angeklickt werden (vgl. Grossman, 2014, S. 2661).

4.3. Auswertung der replizierten Experimente und der unterschiedlichen Situationsräume zur Informationsvermeidung und deren Bedeutung

Insgesamt nehmen an den Experimenten von Grossman 344 Probanden teil, von denen 172 Diktatoren (Spieler X) sind. Im Durchschnitt verdienen die Probanden \$10,53. In diesen Betrag sind die Teilnahmezahlung von \$5 und die bei richtiger Einschätzung gezahlte Auszahlung mit eingerechnet. Dabei verdienen die Diktatoren im Schnitt minimal mehr (\$11,40) als ihre Empfänger (\$9,67). Bei Betrachtung des Default NR, des Active Choice und des Default R Experiments wandelt sich die Entscheidungsquote langsam mit der Art und Weise, wie die Informationsentscheidung abgefragt wird (vgl. Abbildung 4). Während sich bei den Default NR Bedingungen 19 von 42 (45%) für das Nichterhalten von Informationen entscheiden, handeln bei Active Choice mit nur anteilig 25% (10 von 40) signifikant weniger (Standardnormalverteilungsquantil $Z = 1,92$, mit einem Signifikanzniveau von $p = 0,03$) genauso. Im Default R Experiment entscheiden sich fast alle für das Erhalten von Informationen, nur 3% (1 von 19) entscheiden sich dagegen. Dieses extrem niedrige Ignoranz-Level unterscheidet sich signifikant von den Raten, die aus den beiden ersten Experimenten resultieren (mit $Z = 3,85$, $p = 0,001$ bei Default NR und $Z = 2,41$, $p = 0,01$ bei Active Choice). Daraus wird geschlossen, dass die Neigung ignoranten Verhalten zu zeigen (vgl. DWK), was insbesondere in dem Default NR Experiment abgebildet wird, darauf zurück zu führen ist, dass die Variante passiv gewählt werden kann. Auch im Strategy Method Experiment ist die Ignoranzrate insgesamt signifikant niedriger ($Z = 1,77$, $p = 0,04$) als im Default NR Experiment (45%) (vgl. Tabelle 6). Nur 9 der 35 Diktatoren (26%) entscheiden sich dafür, später keine Informationen über die Auszahlung an den Empfänger zu erhalten. Bei Diktatoren, die in beiden Spielvarianten A wählen, ist der Anteil derer, die ignorant sind, mit

29% (5 von 17) nur minimal größer als bei denen, die bei Interessenkonflikt B wählen und bei gleichen Interessen A wählen (hier: 27%; 4 von 15). Aus diesen Ergebnissen des letzten Experiments kann abgeleitet werden, dass Diktatoren, die sich in Situationen, in denen zusätzliche Informationen ihre Wahl beeinträchtigen könnten, lieber ignorant bleiben, trotzdem gerne wissen wollen, wie sich die finale Auszahlung gestaltet. Als nächstes werden die replizierten Experimente von DWK betrachtet. Tabelle 7 zeigt, dass die Ergebnisse aus dem Baseline und dem Default NR Experiment sehr ähnlich zu denen von DWKs Baseline sowie ihrem Hidden Information Experiment sind, was die Robustheit der Ergebnisse bestätigt. Auch Grossman kann einen Unterschied auf einem Signifikanzniveau von $p = 0,02$ ($Z = 2,00$) zwischen den Anteilen der Probanden, die sich für eine unfaire Auszahlung entscheiden, unabhängig davon, ob sie dies wissend oder unwissend tun (Verhaltensweise wird als „inkonsistent“ bezeichnet), feststellen. Sind es im Baseline Experiment nur 9 von 26 (35%), votieren im Default NR Experiment 25% mehr für die Option A (25 von 42, 60%). Bei der Neuauflage der beiden Experimente von Grossman kommt es durch kleine unterschiedliche Maßnahmen zu geringen Unterschieden, die aber allesamt nicht statistisch signifikant sind. So entscheiden sich im Baseline Experiment bei Grossman 35% für ein unfaires Ergebnis und bei DWK nur 26%, was einer Standardabweichung von $Z = 0,59$ entspricht. Bei Default NR beträgt die Standardabweichung $Z = 0,55$ mit 60% gegenüber 53% die A wählen. Keine Informationen wollen bei Grossman 45% bei DWK 44% ($Z = 0,13$). Anteilig von denen, die sich für das Aufdecken vom Spiel entscheiden, votieren bei konträren Interessen 54% bei Grossman und 75% bei DWK für B ($Z = 0,97$). Ignorant bleiben und im Anschluss egoistisch handeln bei der Neuauflage 89% und bei dem Original Experiment 86% ($Z = 0,33$). Das zeigt, dass die Ergebnisse sehr zuverlässig zu replizieren sind. Jemand, der sich in dem Baseline Experiment für Option A (6, 1) entscheidet, handelt inkonsistent zu seinen Fairnessidealen. Deswegen werden in Tabelle 7 und Tabelle 8 alle Probanden, die in dem Baseline, Default NR, Active Choice und Default R Experiment zum Nachteil des Empfängers handeln, unabhängig davon, ob sie dies wissend oder unwissend tun, als inkonsistent tituliert. Wie schon weiter oben angesprochen, ist der Verhaltensunterschied der Probanden vom Baseline Experiment zum Default NR Experiment enorm. Die Ignoranz-Rate wird durch die Art der Angabe der Standardauswahl immer weiter eingedämmt, genauso wie der moralische Spielraum für unpopuläre Entscheidungen. Dies zeigt sich insbesondere bei den Active Choice Bedingungen, unter denen genauso viele Probanden inkonsistent bzw. unfair entscheiden wie im Baseline Experiment, eben exakt 35%. Noch eine geringere Rate weist das Default R Experiment mit 28% auf. Aus der obigen Auswertung ergibt sich, dass Probanden insbesondere dann ihren moralischen Spielraum für egoistisches Verhalten voll ausschöpfen, wenn sie die Möglichkeit dazu haben, Ignoranz passiv zu wählen (vgl. Grossman, 2014, S. 2662f.).

4.4. Kritische Auseinandersetzung und Diskussion

Die Experimente von Grossman unterstreichen nicht nur den Wahrheitsgehalt der Ergebnisse von DWK und somit die Tatsache, dass Probanden gezielt Informationen vermeiden um egoistische Entscheidungen treffen zu können, sondern untersuchen mit ihrer Forschung den Punkt „Informationsvermeidung“ genauer. Situationen, in denen passiv entschieden werden kann, dass keine Informationen folgen sollen, werden von 45% (vgl. Grossman, 2014, S. 2662) gezielt genutzt. Ist die Fragestellung neutral oder muss der Proband sich aktiv dafür entscheiden, nicht informiert zu werden, geht der Anteil der Informationsverweigerer auf ein Viertel bzw. 3% zurück. Dabei spielt die schon zuvor markierte Standardauswahl ebenfalls eine bedeutende Rolle (vgl. Madrian und Shea, 2001, S. 1149; Johnson und Goldstein, 2003, S. 1338). Dies gilt insbesondere in Situationen, die auch eine soziale Komponente aufweisen oder wenn Konsumenten besondere moralische Ideale besitzen, die sie zu wahren versuchen (vgl. Nyborg, 2011, S. 271). Dazu zählen ebenfalls Verteilungsentscheidungen. Die Standardauswahl kann dem Probanden als Vorschlag dienen oder ihm in gewisser Weise als Erlaubnis dienen eine bestimmte Wahl zu treffen. Beim alltäglichen Einkauf ist für die Standardauswahl unser Supermarkt verantwortlich. Anhand der Präsentation seiner Produkte, können Informationen über ihren Herstellungsverlauf mehr oder weniger aktiv in Erfahrung gebracht werden. Insbesondere bei überdurchschnittlich günstigen Produkten oder teuren Bio-Produkten beeinflusst dies das Informationsverhalten und damit die spätere Kaufentscheidung. Beim Beispiel Schokolade entscheiden die Betreiber des Supermarktes und die Schokoladenhersteller maßgebend über die Art der Informationsauswahl. Entsprechend bieten diese, je nach dem unter welchen Bedingungen ihre Produkte bzw. ihre Zutaten hergestellt werden, unterschiedliche Informationsspektren an (vgl. Mistrati, 2010; Drewes, 2012, S. 1f.). Wie gut weiß der Verbraucher Bescheid? Und wie viel wird ihm mehr oder weniger aktiv als „Basiswissen“ zur Verfügung gestellt oder direkt mitgeteilt? Damit ist die Vorauswahl ein Wegweiser des Unterbewusstseins. Auf das Default NR, Active Choice und Default R Experiment passt die Interpretation, dass je aktiver der Agent sich für das Vermeiden von Informationen durch die Standardauswahl entscheiden muss, desto weniger macht er es. Das gilt jedoch nicht für das Strategy Method Experiment. Hier sind bei den Informationsfragen die gleichen Ausgangsbedingungen gegeben wie bei dem Default NR Experiment. Daraus könnte gefolgert werden, dass sich ähnlich viele Probanden gegen Auszahlungsinformationen entscheiden. Das Gegenteil ist jedoch der Fall. 74% wollen informiert werden. Die Informationen erhalten sie jedoch erst nach ihrer Wahl und das ist an dieser Stelle der ausschlaggebende Punkt. Dies liefert einen weiteren Beweis dafür, dass Diktatoren Situationen vermeiden, in denen sie informiert eine Entscheidung treffen müssen, welche ihre Handlungen beeinflussen könnten. Die Strategy Method erweist sich an dieser Stelle als gute Wahl, da hier anders als bei dem Spielvorgehen in den anderen Experimenten die Emotionen der Teilnehmer weniger stark beeinflusst werden und die Methode

für aussagekräftige Ergebnisse steht (vgl. Fischbacher et al., 2012, S. 897; Casari und Cason, 2009, S. 159). Doch welche Motive stecken hinter dem Vermeiden von Informationen? Dafür gibt es unterschiedliche Erklärungsansätze, deren sich auch Grossman bedient. Charness und Jackson, aber auch Chakravarty, Harrison, Haruvy und Rutsröm begründen dieses Verhalten damit, dass die Probanden in Abhängigkeit von ihrem Sinn für Verantwortung moralisch handeln (Charness und Jackson, 2009, S. 246; Chakravarty et al., 2011, S. 901). Ein anderer Ansatz ist, dass durch das Vermeiden von Informationen auch dem eigenen Selbstbild kein Schaden zugefügt werden kann (vgl. Gneezy et al., 2012, S. 7240). Zusätzlich kommt die verhaltenswissenschaftliche Argumentation ins Spiel. Forschungsergebnisse von Krupka und Weber belegen, dass Menschen härter verurteilt werden, wenn sie wissend egoistisch handeln, als wenn sie dieselbe Entscheidung unwissend treffen (vgl. Krupka und Weber, 2013, S. 501). Ob Unwissenheit jedoch von Verantwortung entbindet, wird an dieser Stelle offen gelassen. Eine andere Folgerung berücksichtigt, dass es Individuen wichtig ist, wie sie auf sich und andere wirken. Dabei stellt sich für sie die Frage, ob es imageschädigend ist wissend bzw. unwissend zu handeln. Dieser interessante Erklärungsansatz wird leider nur kurz erwähnt, aber durchaus durch die Tatsache gewichtet, dass auch Bénabou und Tirole das Selbstbildnis als Grund für strategische Ignoranz aufführen (vgl. Bénabou und Tirole, 2006, S. 1652). Aufgegriffen wird er in späteren Forschungen von van der Weele und Grossman (vgl. Grossman und Van der Weele, 2017, S. 173), die im Fortlauf dieser Arbeit untersucht werden. Abschließend ist festzustellen, dass Grossman in seiner Arbeit gleich zweimal beweist, dass Diktatoren gezielt Informationen vermeiden, wenn diese ihre Entscheidung beeinflussen könnten. Zum einen durch die Replikation zweier Experimente von DWK und zum anderen in seinem Experiment Strategy Method, in welchem die Probanden erst nach ihrer Entscheidungsfindung Informationen erhalten können. Einen ganz neuen Blickwinkel eröffnet Grossman durch die Erkenntnis, dass Probanden Informationen im Wesentlichen nur dann vermeiden, wenn die Informationsentscheidung in einer intransparenten Situation passiv wählbar ist.

5. Selbstbild und Ignoranz in gesellschaftlichen Entscheidungen

Das Paper „Self-Image and Willfull Ignorance in social decisions“ von Grossman und van der Weele baut gezielt auf Forschungsergebnissen von früheren selbst verfassten Arbeiten auf, wie die im letzten Kapitel betrachtete „Strategic Ignorance and the Robustness of Social Preferences“ (Grossman, 2014, S. 2659) oder weiteren wie „Self-signaling and Social-Signaling in giving“ (Grossman, 2015, S. 26) und „Inconvenient Truths: Determinants of Strategic Ignorance in Moral Dilemmas“ (Van der Weele, 2014, S. 1). Zuvor gewonnene Erkenntnisse werden miteinander verknüpft und in zusammenhängenden Kontext neu erfasst. Fundamental sind erneut die Forschungsergebnisse von DWK. Grossman und van der Weele versuchen mit ihrer Arbeit gezielt an Forschungen aus

dem Bereich der Sozialpsychologie (vgl. Bandura et al., 1996, S. 2364f.; Sweeny et al., 2010, S. 340f.), die sich mit dem Bayesschen Selbstsignalisierungsmodell beschäftigen, anzuknüpfen. Informationsvermeidung bei hoher Wahrscheinlichkeit, dass eigennütziges Handeln ungünstige Wohlfahrtsfolgen als Konsequenz hätte, ist eine der Hauptursachen für sozial schädliches Verhalten. Ausgehend von dem Bayesschen Modell kreieren van der Weele und Grossman einen Agenten, der sich um sein Image sorgt und die Möglichkeit geboten bekommt mehr über für ihn kostspielige Sozialleistungen zu erfahren. Durch Verschleiern der Präferenzen der Entscheidungsträger können sie zeigen, dass bewusste Ignoranz als Entschuldigung für egoistisches Verhalten dienen kann, und können die Vorstellung von einem Individuum wahren, dass bei vollen Informationen tugendhaft gehandelt hätte. Aus ihrem eigens entworfenen Modell leiten sie verhaltensorientierte Vorhersagen ab, die zum Teil weder mit erfolgsorientierten noch mit sozialorientierten Lösungs- und Erklärungsansätzen vereinbar sind. Diese werden in dem von ihnen entworfenen Experiment getestet. Ihre Ergebnisse sowie die Resultate anderer Experimente untermauern ihre Behauptungen und können damit die Theorie von der Bedeutung der Imagewirkung (vgl. Ariely et al., 2009, S. 544) erweitern.

5.1. Vermeidung von Informationen über nachteilige Wohlfahrtsfolgen

Die Arbeit „Self-Image and Willful Ignorance in Social decisions“ von Grossman und van der Weele ist das Ergebnis über zehn Jahre langer Forschung auf dem Themengebiet „Eigennütziges Verhalten und bewusste Ignoranz von Informationen in Entscheidungssituationen“. In ihrer Veröffentlichung greifen sie auf ihre vorherigen Resultate zurück und machen sich verhaltenspsychologische Erkenntnisse zu Nutze. Durch diese Herangehensweise kann das Entscheidungsverhalten auf einer neuen Ebene erforscht werden und damit auch die Forschungsfragen dieser Arbeit in einer zuvor nicht möglichen Tiefe beantworten. Dabei spielt die Selbstwahrnehmung einer Person eine entscheidende Rolle. Individuen werden gerne als ehrlich und altruistisch wahrgenommen (vgl. Mazar et al., 2008, S. 639f.). Sie haben regelrecht Angst davor, als schlechter Mensch verurteilt zu werden (vgl. Norgaard, 2006, S. 348f.). Langezeit war dies ein Themenspektrum, das Psychologen vorbehalten war und von der Klimaforschung aufgegriffen wurde (vgl. Bateman und O'Connor, 2016, S. 214; Grossman und Van der Weele, 2017, S. 174). Bei aktueller Thematik vermeiden die Menschen es bewusst, über die Folgen des Klimawandels zu erfahren, aus Angst den eigenen daraus resultierenden Ansprüchen nicht zu entsprechen und Selbstzweifeln ausgesetzt zu sein. Trotz der Logik der eben aufgestellten Argumentationskette stellt sich jedoch die Frage: Wie kann Ignoranz überhaupt als entlastende Lösung für das Gewissensproblem dienen, wenn sich in gewisser Weise aktiv dafür entschieden wird ignorant zu sein? Durch das Aufstellen von Verhaltensmustern kann dieses Paradoxon erklärt werden und ein tiefergehendes Verständnis für bewusste Ignoranz in diese Arbeit eingebracht werden.

5.2. Einführung in das Modell

Im Kontext des Bayesschen Selbstsignalisierungsmodells stellen Grossman und van der Weele für ihr Vorgehen drei Thesen auf, die sie im Fortlauf ihrer Arbeit beweisen möchten. Als erstes möchten sie zeigen, dass ein Gleichgewicht existiert, in dem weniger altruistische Personen, welche um ihr Image besorgt sind, es strikt bevorzugen ignorant zu sein. Darauf folgend ist ihre zweite These, dass aus dem Gleichgewicht fünf unterschiedliche, beweisbare Verhaltensmuster resultieren, die sie laut ihrer dritten These, in dem von ihnen entworfenen Experiment einheitlich beweisen wollen. Zu dem Experiment soll ein Modell hinführen, das im Anschluss der Beschreibung der Verhaltensmuster vorgestellt wird (vgl. Grossman und Van der Weele, 2017, S. 174f.). Das erste Verhaltensmuster besagt, dass wenn zu Beginn Unsicherheit über den sozialen Nutzen einer Handlung besteht, selbst wenn es Möglichkeiten gibt, diese Intransparenz zu entschleiern, dies tendenziell die Anreize, sich prosozial einzubringen, schwächt. Das zweite Verhaltensmuster sagt aus, dass sich gerade altruistische Typen mit niedrigen prosozialen Kosten informieren. Dies impliziert unter anderem, dass Individuen, die sich aktiv dazu entscheiden, Informationen erhalten zu wollen, sich automatisch sozialer verhalten als Personen, die dieselben Informationen passiv erhalten haben. Dieses Phänomen wird als Einsortierung bezeichnet (vgl. Grossman und Van der Weele, 2017, S. 175). Das nächste Verhaltensmuster titelt mit der Behauptung „Vorsätzliche Ignoranz entschuldigt“. Jemand, der vollständig informiert handelt und sich im Klaren über die Konsequenzen seiner Entscheidungen ist, wird viel härter verurteilt als jemand, der dieselben Entscheidungen uninformiert trifft. Selbst dann wenn er sich aktiv für die Ignoranz von Informationen entschieden hat. Hat der Entscheidungsträger erst einmal seine altruistischen Vorlieben preisgegeben, so ist es schwer diesen Charakterzug weiter mit Ignoranz zu verschleiern, lautet die These des vierten Verhaltensmusters. Diese Situation tritt insbesondere dann auf, wenn die Information die Entscheidungssituation nicht mehr beeinträchtigen kann (vgl. Strategy Method). Das fünfte Verhaltensmuster verrät, dass Agenten, die sich im Gleichgewicht für Ignoranz entscheiden, sich ohne Informationen strikt besser stellen (vgl. Grossman und Van der Weele, 2017, S. 176). Das Modell, das Grossman und van der Weele entwerfen, lehnt sich an das Modell von Bénabou und Tirole (2006, S. 1652f.) an, die ein vergleichbares Modell verwenden, um bewusste Ignoranz im Kontext von Tabus zu analysieren, und ähnelt auch dem von Grossman (2013, S. 26f.), der sich mit der Selbst und Fremdwahrnehmung in Verteilungssituationen beschäftigt und dabei Vorhersagen des Bayesschen Selbstsignalisierungsmodell in seinem wahrscheinlichkeitstheoretischen Diktator-Spiel testet. Warum entscheiden sich Personen nichts über die Konsequenzen ihres Handelns erfahren zu wollen? Dazu führen Grossman und van der Weele ein Präferenzen-Signalisierungsmodell ein, das die intrinsische Sorge für die soziale Fürsorge mit der Vorliebe als prosozial Handelnder wahrgenommen zu werden kombiniert und damit persönliche Präferenzen materiellen Auszahlungen vorzieht. Eine Interpretationsmög-

lichkeit für diese Art von Modellen ist die „Selbstsignalisierende Interpretation“ mit einem internen Betrachter. Die Person ist in der Vorstellung innerlich in zwei Hälften gespalten. Es gibt das informierte Entscheidungstreffende-Selbst, das über seine Vorlieben Bescheid weiß und dem entsprechend auch handelt. Sein Ziel ist es das andere, Beobachter-Selbst, das keinerlei Informationen hat und dem es an Wissen über seine Vorlieben mangelt und deswegen auch als Freudsches „Super-Ego“ zu interpretieren ist, zu beeindrucken (vgl. Grossman und Van der Weele, 2017, S. 177). In ihrem Modell entwerfen die beiden Verfasser auf Grundlage dieses Präferenz-Signalisierungs-Modells einen Agenten. Dieser entscheidet zunächst, ob er sich an einer prosozialen Handlung beteiligt oder nicht. Wenn er sich dafür entscheidet prosozial zu handeln, fallen für ihn materielle Kosten an und es ergibt sich ein ungewisser Wohlfahrtseffekt. Dass ein bestimmter Wohlfahrtsfaktor eintrifft, kann mit einer bestimmten Wahrscheinlichkeit vorher gesagt werden. Der komplementäre Wahrscheinlichkeitsanteil bedeutet, dass es zu keiner Wohlfahrtssteigerung kommt, wobei der Wohlfahrtsfaktor größer als die materiellen Kosten sind. Bevor der Agent eine Entscheidung trifft, kann er sich überlegen, ob er informiert werden will, wobei diese Informationen mit einem Informationskostenfaktor verbunden sind. Wenn er sich informieren lässt, kann er erfahren, dass die Wohlfahrt einem bestimmten Wert entspricht, jedoch kann sie auch Null sein. Für den entworfenen Agenten wird eine Nutzenfunktion, die jedoch nur dem Entscheidungstreffenden-Selbst bekannt ist, aufgestellt. Diese beinhaltet einerseits einen Gewichtungsfaktor, der seine sozialen Präferenzen bzw. den Grad seiner prosozialen Motivation widerspiegelt, und andererseits einen zweiten Gewichtungsfaktor, der die Vorteile eines positiven Selbstimages darstellen soll. Die Nutzenfunktion des Agenten lässt sich in drei Bausteine gliedern. Der erste Teil ist der materielle Nutzen, der jedoch nur dann zum Tragen kommt, wenn der Agent sich dazu entschließt prosozial zu handeln. Wenn dem so ist, ergibt sich in diesem Teil die Differenz aus dem Grad der prosozialen Motivation multipliziert mit dem erwarteten Wohlfahrtsnutzen und den materiellen Kosten. Der erwartete Wohlfahrtsnutzen kann einem bestimmten Wert oder Null entsprechen, wenn Informationen eingeholt wurden, oder das Produkt aus der Wahrscheinlichkeit und einen bestimmten Wohlfahrtswert sein, wenn keine Informationen vorliegen. Im zweiten Baustein sind die Informationskosten zu subtrahieren und im dritten Teil wird der Faktor, der die Vorteile eines positiven Selbstimages darstellt, multipliziert mit den nachträglichen Erwartungen des Beobachter-Selbst, bezogen auf den sozialen Präferenzenparameter, addiert. Obwohl alle Informationen dazu dienen den Charakter des Agenten genauer zu identifizieren, sind sie nicht wirklich nachweisbar. Das heißt, dass das Beobachter-Selbst, das einen Mangel an Informationen hat, am besten damit beraten ist, seine Einschätzungen von dem Faktor des Grads prosozialer Motivation und dem Vorteilsfaktor für ein prosoziales Selbstimage, welche vom Informationsstand abhängig sind, stets zu aktualisieren. Dadurch, dass das Beobachter-Selbst dem Entscheidungstreffenden-

Selbst Vertrauen schenkt, werden die Beobachtungen des Beobachtenden-Selbst beeinflusst. In Bezug auf Heterogenität trifft das Modell zwei Annahmen. Die erste ist, dass mit einer gegen Null gehenden Wahrscheinlichkeit davon ausgegangen werden kann, dass es sich bei dem Agenten um einen Anti-Sozialen Agenten bzw. einen Homo oeconomicus handelt. Dieser kümmert sich nur um seinen persönlichen materiellen Nutzen und hat bei den Faktoren der sozialen Präferenzen sowie dem Vorteilsfaktor in Bezug auf ein soziales Image, die Angabe Null stehen (vgl. Grossman und Van der Weele, 2017, S. 179). Die zweite Annahme ist, dass mit der entsprechenden Gegenwahrscheinlichkeit zur ersten Annahme es sich um einen sozialen Agenten handelt, der sich um sein soziales Image und die soziale Wohlfahrt sorgt. Dabei soll der Vorteilsfaktor eines prosozialen Images für alle Agenten dieser Art gleich sein, einerseits größer als Null und andererseits kleiner als die Kosten einer prosozialen Handlung. Dies schließt eine prosoziale Handlung sowohl aus Imagegründen als auch Situationen, aus denen ein aushaltbares Gleichgewicht resultiert, aus (vgl. Grossman und Van der Weele, 2017, S. 180). Der Zeitplan für das Spiel gestaltet sich wie folgt: Zunächst entscheidet sich per Zufall, wie die Wohlfahrt ausfällt, die prosoziale Handlung sowie die Typisierung des Agenten. Danach entscheidet sich der Agent, ob er informiert werden möchte oder ein leeres Signal empfangen möchte und erhält dieses. Danach legt der Agent fest, ob er prosozial handeln möchte. Am Schluss nimmt das Beobachter-Selbst die Handlung sowie die Informationen wahr und aktualisiert mit Hilfe dieser Informationen seine bisherige Einschätzung über das Entscheidungstreffende-Selbst. Aus der Abfolge des Spiels ergeben sich verschiedene Unsicherheiten. Zum einen ist das Beobachter-Selbst im Unklaren über den wahren Charakter des Agenten, zum anderen sind beide Selbst zu Beginn unsicher über den Betrag der Wohlfahrt. Ob diese Intransparenz aufgedeckt wird, liegt allein in der Entscheidung des Entscheidungstreffenden-Selbst (vgl. Grossman und Van der Weele, 2017, S. 180).

5.2.1. Lösungskonzept

Grossman und van der Weele treffen unter Verwendung des perfekten bayesschen Gleichgewichts die Vorhersage, dass alle Typen eine nutzenmaximierende Strategie haben. Die Überlegung, dass der Erwartungswert in Abhängigkeit von dem sozialen Präferenzfaktor und der prosozialen Handlungsentscheidung mit Informationsmöglichkeit, dem Erwartungswert in Abhängigkeit von dem sozialen Präferenzfaktor und der prosozialen Handlungsentscheidung mit Informationsmöglichkeit sowie der Maximierungsstrategie, die hier dem Gleichgewichtsstrategieprofil entspricht, die aus der Adaption der Bayesschen Regel folgt. Sie stellen die Regel auf, dass Agenten, die indifferent sind, ob sie Informationen annehmen wollen, sich dafür entscheiden. Damit soll bewusste Ignoranz konservativ bewiesen werden (vgl. Grossman und Van der Weele, 2017, S. 180). Generell liegt der Fokus auf semi-separativen Gleichgewichten, bei denen stets ein Grenzwert für den Grad prosozialer Motivation existiert. Ist der Grad prosozialer Motivation bei einem Agenten höher

als der Grenzwert im Gleichgewicht, entscheidet sich der Typ für die Informationsakquise, und wenn nicht dagegen. Sobald das Beobachter-Selbst seinen Standpunkt festgelegt hat, kann ein stabiles Gleichgewicht garantiert werden. Im Falle eines Ignoranz-Gleichgewichts muss der Agent zwischen den erwarteten Kosten und der zu erwartenden Selbstanerkennung indifferent sein (vgl. Grossman und Van der Weele, 2017, S. 180f.).

5.2.2. Existenz eines bewussten Gleichgewichts bei Ignoranz

Das wichtigste theoretische Ergebnis der Arbeit von Grossman und van der Weele ist die Existenz eines Gleichgewichts, in dem es soziale Agenten bevorzugen ignorant zu bleiben. Damit stellen sie ihre erste These im Rahmen ihres Modells auf.

1. These: Es existiert eine Grenzwahrscheinlichkeit kleiner als eins und Informationskosten, die ihre Grenze bei einem Wert größer oder kleiner als Null haben, so dass, wenn die Wahrscheinlichkeit größer als die Grenzwahrscheinlichkeit ist und die Informationskosten im Intervall zwischen den beiden Grenzen liegen, ein semi-separatives Gleichgewicht vorliegt, das durch den Grad prosozialer Motivation charakterisiert wird. So entscheidet sich der Homo oeconomicus erstens gegen eine prosoziale Handlung und akquiriert nur dann Informationen, wenn diese kleiner oder gleich Null sind. Zweitens bleiben alle sozialen Typen mit einem niedrigeren Grad prosozialer Motivation als der Gleichgewichtsgrenzwert ignorant und führen keine prosoziale Handlung aus, während alle Typen, mit einem höheren Grad prosozialer Motivation Informationen einfordern und prosozial handeln, jedoch nur wenn die Information über die Wohlfahrtssteigerung ihrer Handlung einem hohen Wert entspricht.

Um nun zu verstehen, warum soziale Typen ignorant bleiben, selbst wenn dies kostspielig ist, muss man die Austauschbeziehungen bei der Informationsentscheidung betrachten. Im Gleichgewicht sind die materiellen Vorteile für den Entscheidungsträger und das Image, das mit der Unwissenheit einhergeht, sicher. Bewusste Ignoranz ist stets gefolgt von dem Vermeiden prosozialen Verhaltens und den damit verbundenen Ausgaben. Jedoch ziehen sie negative Enttäuschungen, die sich aus den eigenen prosozialen Präferenzen ergeben, nach sich. Dieses Verhalten wirkt sich auch unvorteilhaft auf das Image, das nicht mehr mit den zuvor gehegten Erwartungen einhergeht, aus. Zu Beginn sind sowohl die materielle Auszahlung als auch das Image, das aus der Informationsakquise resultiert, ungewiss. Wenn sich später herausstellt, dass prosoziales Verhalten keinerlei Wohlfahrtssteigerung hervorruft, kann sich der Agent in die Gruppe der altruistischen Agenten einreihen, ohne nennbare Verlust gemacht zu haben. Sollte sich aber zeigen, dass prosoziales Agieren Wirtschaftsvorteile bringt, sitzt der Agent

in der Zwickmühle. Entscheidet er sich, die Kosten für prosoziales Handeln zu tragen, erfährt er das höchstmögliche Erscheinungsbild. Entscheidet er sich hingegen für das Gegenteil, erhält er die negativste Imagebewertung. Diese besitzt sonst nur der Homo oeconomicus, der selbst dann eine prosoziale Entscheidung ablehnt, wenn sie ihn nicht einmal etwas kosten würde. Das heißt, dass sich Agenten, denen ihr Selbstimage grundsätzlich wichtig ist, aber niedrigere soziale Präferenzen haben als der Gleichgewichtsgrenzwert, für Informationsvermeidung entscheiden, wenn die Wahrscheinlichkeit hoch ist, dass prosoziales Handeln einen Wohlfahrtsvorteil verursacht. Das bedeutet, dass Ignoranz im Kontext dieses Gleichgewichts dazu dient in Situationen, die soziales Verhalten einfordern, die wahren Absichten, die der Agenten hat, zu verschleiern. Demzufolge schützt Ignoranz auch das Selbstbild vor dem scharfen Urteil des Beobachter-Selbst. Dieses kann nämlich nicht widerlegen, dass der vollständig informierte Agent ausreichend altruistisch ist, so dass er prosozial gehandelt hätte (vgl. Grossman und Van der Weele, 2017, S. 182).

5.3. Empirische Implikationen des Modells

In diesem Abschnitt werden, aufbauend auf das zuvor aufgestellte Gleichgewicht der ersten These des Modells, empirische Implikationen, die sich jeweils mit dem entsprechenden Verhaltensmuster verknüpfen lassen (vgl. Kapitel 5.2), abgeleitet (vgl. Grossman und Van der Weele, 2017, S. 184).

5.3.1. Informationsvermeidung

Unter Berücksichtigung der Experimente zur bewussten Ignoranz von DWK soll nun das Selbstsignalisierungsmodell getestet werden. Die erste Frage, die sich stellt, ist, ob es die eben genannten Ergebnisse erklären kann, die mit den Theorien der auszahlungsorientierten Präferenzen und der sozialen Signalisierung unvereinbar sind. Um dies zu untersuchen, werden zwei verschiedene Experimente unterschieden, zum einen das Entscheidungsspiel (vgl. Zeitplan Kapitel 5.2). Hier entscheidet der Teilnehmer zunächst, ob er informiert werden möchte oder nicht und im Anschluss, ob er prosozial handelt, was dem Vorgehen des Hidden Information Experiments entspricht. Zum anderen sinnbildlich dem Baseline Experiment entsprechend, ein Experiment, in dem allgemein bekannt ist, dass eine prosoziale Handlung eine bestimmte Wohlfahrt hervorruft und der Entscheider sich nur für diese oder gegen diese entscheiden muss. Um den Anteil der Agenten, die sich in im transparenten Spiel egoistisch verhalten, mit dem Anteil zu vergleichen, die im Entscheidungsspiel unwissend bleiben wollen, wird zunächst das Gleichgewicht für das transparente Spiel abgeleitet. Wie schon im Entscheidungsspiel in Kapitel 5.2.2 wird der Fokus auf ein semiseparierendes Gleichgewicht, das sich durch einen Grenztyp für prosoziale Motivation auszeichnet, gelegt. Es wird davon ausgegangen, dass für dieses Spiel eine Stabilitätsbedingung existiert, damit die Einzigartigkeit des Grenztyps gewährleistet werden kann. Dazu wird der Faktor der Imageanerkennung mit einbezogen. Der Schwellentyp ist genau der, dessen

Nettokosten der prosozialen Aktion den Imagevorteilen entsprechen. Grossman und van der Weele vergleichen dieses Gleichgewicht mit dem Gleichgewicht bei Ignoranz im Entscheidungsspiel, wobei sie, um mit den Experimenten in der Literatur konsistent zu sein die Annahme treffen, dass die Informationskosten kostenlos sind. Dabei können sie mit der folgenden These abschließen.

2. These: Es existiert ein Mittelwert für die Vorteile eines positiven Selbstimages sowie für die Wahrscheinlichkeit, die kleiner als eins ist. Außerdem ist der gegen Null gehende Anteil an Homo oeconomicus kleiner als ein halb, sodass, wenn die Vorteile eines positiven Selbstimages, sowie die Wahrscheinlichkeit über ihren und die gegen Null gehenden Anteile für einen Homo oeconomicus unter ihren Mittelwerten liegen, dann der Anteil aller Agenten, sowohl der sozialen Agenten als auch der Homo oeconomicus Agenten, die sich für Ignoranz im Gleichgewicht entschieden haben, im Entscheidungsspiel höher als der Anteil im transparenten Spiel, die dort eigennützig entscheiden, ist (vgl. Grossman und Van der Weele, 2017, S. 184).

Diese These zeigt, dass das Signalisierungsmodell die Vermeidung von Informationen unter der richtigen Konstellation von Parametern erklären kann. Zur Informationsvermeidung kommt es, wenn die sozialen Typen es umgehen wollen dasselbe Sozialimage wie der Homo Oeconomicus zu erhalten. Dies hat unterschiedliche Auswirkungen auf ihr Verhalten in den beiden Spielen. Im transparenten Spiel agiert der Homo Oeconomicus nicht sozial. Dies erhöht den Signalwert einer prosozialen Aktion und veranlasst einige marginal soziale Typen sich prosozial zu verhalten. Im Gegensatz dazu möchte der Homo Oeconomicus im Entscheidungsspiel Informationen erhalten. Dies verringert den Imagewert für die Informationsbeschaffung und verursacht, dass sich schwach soziale Typen auf einmal für Ignoranz entscheiden und sich die Anzahl egoistischer Entscheidungen im Anschluss erhöht. Die Bedingungen für die Parameter sind in diesem Fall ausreichend um sicher zu stellen, dass die Menge der sozialen Typen, die sich für Unwissenheit im Entscheidungsspiel umentscheiden, die der Homo Oeconomicus ausgleicht, dadurch dass er sich informiert. Wichtig ist, dass erstens der Anteil der Homo oeconomicus nicht zu groß ist, da sonst anteilig zu geringe Unwissenheit herrscht. Zweitens muss die Bedeutung des Imageanliegens ausreichend hoch sein, damit genügend soziale Typen dazu veranlasst werden Intransparenz vor zu ziehen. Und drittens sollte die Wahrscheinlichkeit, dass ein prosoziales Handeln einen positiven sozialen Nutzen auslöst, hoch genug sein, um Unwissenheit für die sozialen Typen im Entscheidungsspiel attraktiv zu machen. Wenn diese Bedingungen nicht erfüllt sind, kann es trotzdem zu vorsätzlicher Ignoranz kommen, aber der Anteil der ignoranten Agenten gleicht je nach dem nicht die aus, die sich im transparenten Spiel egoistisch verhalten. Zusammenfassend lässt sich

sagen, dass Entscheidungen unter vollständiger Information ein klares Signal dafür geben, bei wem es sich um einen egoistischen Typen handelt. Dies ist nicht unbedingt der Fall, wenn die Informationsbeschaffungsentscheidung zwiespältig ist und das Beschaffen von Informationen, das Hinzugruppieren zu nicht-sozialen Typen bedeutet. Diese Verwässerung der Signalisierungsanreize kann erklären, warum vorsätzliche Ignoranz egoistisches Verhalten in Umgebungen ohne Unsicherheit übersteigen kann (vgl. Grossman und Van der Weele, 2017, S. 184f.).

5.3.2. Einsortierung

Im Gleichgewicht entscheiden sich bestimmte Typen für Transparenz. Infolgedessen verhalten sich Menschen, die sich aktiv für die Beschaffung von Informationen entschieden haben, im Durchschnitt anders als diejenigen, die dieselben Informationen passiv erhalten haben. Um diese Idee formal zu untersuchen, wird das Verhalten derjenigen, die sich im Entscheidungsspiel informieren, mit dem Verhalten der Spieler im transparenten Spiel verglichen.

3. These: Sind die Informationskosten größer als Null, ist der Anteil der Agenten, die im transparenten Spiel prosozial agieren, niedriger als der entsprechende Anteil der Agenten, die sich im Gleichgewichtsspiel des Entscheidungsspiels informieren und wissen, dass es zu einem Wohlfahrtsvorteil im Fall einer prosozialen Aktion kommt. Wenn die Informationskosten größer gleich Null sind, gibt es einen Näherungsanteil von Agenten des Typs Homo oeconomicus. Auch wenn dieser kleiner ausfallen sollte, ändert sich nichts am Ergebnis.

Die Intension dieses Ergebnis ist, dass das durchschnittliche prosoziale Verhalten im transparenten Spiel von allen Agenten übernommen wird, wobei der durchschnittliche Altruismus gleich dem erwarteten Wert für soziale Präferenzen ist. Innerhalb dieser Stichprobe verhalten sich nur die altruistischsten Typen prosozial, während es bei den Agenten am Scheidepunkt von den Parametern des Wohlfahrtszuwachses, der Kosten für prosoziales Handeln und den Vorteilen eines positiven Images abhängt. Im Vergleich dazu informieren sich im Entscheidungsspiel meist die altruistischen Typen. Bedingt durch die Informationsbeschaffung ist der Altruismus höher als der Erwartungswert und auch das prosoziale Handeln selbst ist häufiger als im transparenten Spiel. Der Agent vom Typ Homo oeconomicus entscheidet sich für Ignoranz, wenn die Informationsbeschaffung etwas kostet. Deshalb entschließen sich alle anderen Agenten für den prosozialen Einsatz, wenn dieser einen Wohlfahrtsvorteil bringt. Besteht die Möglichkeit, dass Informationen kostenlos sind, entscheidet sich der Agent vom Typ Homo oeconomicus für Transparenz, was die relative Häufigkeit von prosozialem Handeln unter den Wissenden senkt. Deswegen ist es in diesem Fall besonders wichtig, dass die Anzahl der Agenten vom Typ Homo oeconomicus einen bestimmten Grenzwert nicht

überschreitet, da sonst möglicherweise der Anteil von sozial Handelnden unter den Informierten tatsächlich sinkt (vgl. Grossman und Van der Weele, 2017, S. 185f.).

5.3.3. Rechtfertigung

Wenn Typen verschiedenen Aktionen zugeordnet werden, impliziert dies, dass im Gleichgewicht die Handlung, die mit einem bestimmten Image einhergeht, von den entsprechenden Typen ausgeführt wird. Ein egoistisches Image wird hierbei auf Grund der vorhergehenden Informationsentscheidung verliehen. Agenten im Gleichgewicht, die sich im Bewusstsein der negativen sozialen Konsequenzen gegen eine prosoziale Handlung entscheiden, werden als egoistischer verurteilt als Agenten, die dieselbe Entscheidung in Unwissenheit treffen. Liegt der Grad für soziale Präferenzen bei Agenten unter dem Gleichgewichtsgrenzwert, entscheidet sich der Agent stets für Unwissenheit.

1. Logische Folgerung: Agenten, die sich gegen eine soziale Handlung entscheiden, haben ein besseres Image, wenn sie davor ebenfalls Informationen abgelehnt haben, als wenn sie es in dem Wissen einer möglichen Wohlfahrtssteigerung tun.

Diese Folgerung resultiert direkt aus dem Glaubensansatz für das Gleichgewicht, in dem der Erwartungswert der sozialen Präferenzen für eine wissentlich wohlfahrtsbringende soziale Handlung vor dem unwissenden Aussetzen einer Sozialtat und weit vor dem wissentlichen Aussetzen eines wohlfahrtsbringenden Beitrags rangiert. Das bedeutet, dass bewusste Ignoranz tatsächlich als Entschuldigung für sich selbst oder für andere dienen kann, auch wenn klar ist, dass man die tatsächlichen Folgen hätte kennen können (vgl. Grossman und Van der Weele, 2017, S. 186f.).

5.3.4. Bewusste versus unbewusste Ignoranz

Im Gleichgewicht bleibt dem Agenten die Wahl zwischen einer kostspieligen Aktion oder dem Entblößen seiner egoistischen Züge. Daher neigt er automatisch dazu lieber unwissend zu bleiben. Hat der Agent bereits seinen Typ offenbart und steht nach dieser Enthüllung vor der Entscheidung weitere Informationen über den Zustand zu erhalten (vgl. Strategy Method), so kann der Entscheidungsträger nicht länger ignorant bleiben, um seine Entscheidungen, die er unter vollen Informationen treffen würde, zu verdecken. Deswegen wird erwartet, dass mehr Individuen sich für Informationen entscheiden. Um diesen intuitiven Anspruch formell zu begründen, wird ein modifiziertes Spiel betrachtet, welches im Folgenden beschrieben wird. Das Spiel, das unbewusste Ignoranz darstellen soll, zeichnet sich dadurch aus, dass der Zeitpunkt von Informationsentscheidung und Handlungsentscheidung, wie sie im Entscheidungsspiel aufeinander abfliegen, vertauscht ist. Mit Hilfe dieser Regel wird der Agent dazu gezwungen, sein Verhalten vorab preis zu geben, wobei er sich nur zwischen der allgemeinen Wohlfahrt und seiner eigenen entscheiden kann.

4.These: Sind die Informationskosten größer oder gleich Null, ist der Anteil der Agenten, die sich im modifizierten Spiel für Ignoranz entscheiden, niedriger als der Anteil im Gleichgewicht beim Entscheidungsspiel (vgl. Grossman und Van der Weele, 2017, S. 187).

Um die Logik hinter dieser These zu verstehen, hilft es beide Spiele nacheinander zu betrachten. Im Entscheidungsspiel gibt es einen bestimmten Anteil imagebedachter Agenten, ohne Agenten des Typs Homo oeconomicus mit ein zu beziehen, die es bei Informationskosten größer gleich Null strikt bevorzugen keine Informationen zu erhalten. Diese Agenten verzichten auf den instrumentellen Wert von Informationen im Austausch gegen den Imagewert von Ignoranz, wenn sie dies vor der Entscheidung über soziales Handeln tun können. Folglich entsprechend der Regel, dass gleichgültige Agenten aufdecken, wird ein Bruchteil im Entscheidungsspiel aufdecken. Im Gegensatz dazu haben im Spiel für unbewusste Ignoranz weder Informationen noch Ignoranz irgendeinen Wert, nachdem die prosoziale Entscheidung längst getroffen wurde und den Charakter des Agenten aufgedeckt hat (vgl. Grossman und Van der Weele, 2017, S. 187).

5.3.5. Bezahlung für Ignoranz

Das Gleichgewicht kann dazu verwendet werden, vergleichende Statistiken über den Preis der Information abzuleiten. Alle anderen Faktoren gleich behaltend, sorgen positive Informationskosten, also Informationen deren Preis bei kleiner als Null liegt, das heißt, der Agent bekommt Geld, wenn er Informationen annimmt, dafür, dass Unwissenheit für Agenten immer weniger attraktiv wird und der Grad für soziale Motivation niedriger ist, als wenn die Informationen kostenlos wären. Nichts desto trotz stellen sich die sozialen Typen, die eine niedrigere soziale Präferenz als den Grenzwert besitzen, strikt besser, wenn sie für Ignoranz zahlen. Dies ist für sie die beste Alternative, um nicht zusammen mit dem informierten Homo oeconomicus oder mit den ganz altruistischen Typen eingruppiert zu werden und für prosoziales Verhalten bezahlen zu müssen (vgl. Grossman und Van der Weele, 2017, S. 188).

2. Logische Folgerung: Es gibt einen positiven Anteil im Gleichgewicht, der dafür bereit ist, zu bezahlen, um ignorant bleiben zu können (vgl. Grossman und Van der Weele, 2017, S. 188).

5.4. Experimentaufbau und Auswertung zur bewussten Ignoranz

Vorherige Experimente lieferten stets nur den Beweis für eines der Verhaltensmuster, die im Kapitel 5.2 erklärt wurden und im vorherigen Kapitel der empirischen Implikationen erneut aufgegriffen wurden. Auch Grossman und van der Weele machen sich in ihrer gemeinsamen Arbeit erneut die Experimente von DWK zu Nutze (vgl. Kapitel 3) und übernehmen das Experimentdesign zu großen Teilen. So wurde

auch die Einführung für das Experiment erneut übernommen (vgl. Experimentbeschreibung Kapitel 3.2) (vgl. Grossman und Van der Weele, 2017, S. 188f.).

5.4.1. Informationsvermeidung im Experiment

Als erstes soll das Experiment für Informationsvermeidung repliziert werden, wobei damit nicht die Robustheit vorheriger Experimente nachgewiesen werden soll (vgl. Feiler, 2014, S. 256f.; Grossman, 2014, S. 2660; Larson und Capra, 2009, S. 468), sondern es soll einerseits eine Grundlage für Vergleichbarkeit geschaffen werden, andererseits Beweismittel für alle fünf Verhaltensmuster, die aus dem Gleichgewicht implizieren, einen einheitlichen Rahmen finden und zu guter Letzt eigene Ergebnisse für das Hidden Information Experiment von DWK (vgl. Kapitel 3.2.2) aufführen. Zu diesem Zweck führen sie zwei unterschiedliche Experimenthandlungen durch. Zum einen das CIG Only Experiment, das DWK's Baseline Experiment repliziert, in welchem die Probanden wissentlich die Spielvariante mit konträren Interessen spielen, zum andern das Hidden Information Experiment. Dabei steht die Nullhypothese dafür, dass die Rate, mit der die Teilnehmer selbstsüchtig im CIG Only Experiment wählen, mindestens so hoch ist wie die Ignoranzrate im Hidden Information Experiment, während die Alternativhypothese lautet, dass die Ignoranzrate tatsächlich höher ist. Die Ergebnisse in Abbildung 5 zeigen, dass im CIG Only Experiment 9 von 26 (35%) Diktatoren A wählen und im Hidden Information Experiment 72 von 120 (60%) Unwissenheit bevorzugen. Dieser Unterschied ist bei 5% signifikant, bei Anwendung eines einseitigen (oder zweiseitigen) exakten Tests nach Fisher (Fisher-Yates-Test, exakter Chi-Quadrat-Test). Somit sind die Hauptergebnisse von DWK nachgewiesen (vgl. Grossman und Van der Weele, 2017, S. 189f.).

5.4.2. Einsortierung im Experiment

Aufgrund ihrer Aktionen werden Typen zwei verschiedenen Gruppen zugeteilt. Bei der ersten Variante werden Personen, die denselben Kenntnisstand haben – das Wissen, dass sie sich in der Spielvariante von konträren Interessen befinden – aber diesen auf unterschiedliche Weise erreicht haben, entweder bereits von außen gegeben (CIG Only Experiment) oder durch eine positive Informationsentscheidung (Hidden Information Experiment), verglichen. Die zweite Variante stellt getrennt Maßnahmen für den sozialen Typ und den um sein imagebesorgten Typ bereit und vergleicht diese zwischen den Teilnehmern, die unterschiedliche Verhaltensweisen in dem Hidden Information Experiment zeigen. Der erste Ansatz basiert auf der Vorgehensweise von DWK, während der zweite komplett neuartig ist. Zunächst wird die erste Einteilungsvariante betrachtet. Um die Hypothese zu testen, die in der dritten These abgeleitet wurde, wird die Häufigkeit von prosozialem Verhalten im Spiel CIG Only Experiment mit der Häufigkeit, mit der sich Teilnehmer im Hidden Information Experiment für Transparenz entscheiden und sich dann ebenfalls im konträren Interessenspiel wiederfinden, verglichen. Die hellgrauen Balken in Abbildung 5

zeigen den Prozentsatz der Probanden, die unter vollen Informationen eigennützig in der konträren Interessensituation entscheiden. Wenn die Informationen im Hidden Information Experiment endogen sind, wählen 17% der Diktatoren (4 von 24), die sich bewusst im konträren Interessenspiel befinden, selbstsüchtig, wesentlich weniger als in dem CIG Only Experiment, in dem die Informationen exogen gegeben sind und 35% der Diktatoren (9 von 26) egoistisch handeln ($p = 0,13$, FET und $p = 0,084$ im einseitigen z-Test). Die Ergebnisse zeigen, dass der Vergleich des Verhaltens über unterschiedliche Verhaltensgruppen hinweg nicht hoch signifikant ist, aber durchaus konsistent mit den Verhaltensweisen im Gleichgewicht. Nun zur zweiten Einteilungsvariante: Eine alternative Methode zur Untersuchung der Sortierung ist das zusätzliche Erforschen von Diktatortypen und das Vergleichen der Messergebnisse mit den Erklärungsansätzen, die aus dem Gleichgewicht resultieren. Zu diesem Zweck haben Grossman und van der Weele in neun der Hidden Information Experimente zwei individuelle Attributmaße von den Teilnehmern erhoben, nachdem sie das Hidden Information Experiment abgeschlossen und Rückmeldungen über ihre Einnahmen aus diesem Spiel erhalten hatten. Die erste Messung beinhaltete die Bewertung der sozialen Werteorientierung, die mit der Slider-Methode von Murphy et al. (2011, S. 772f.) durchgeführt wurde. In der Aufgabe für die soziale Werteorientierung wird ein Subjekt gebeten, eine Geldverteilung zwischen sich und einem anderen Subjekt in sechs verschiedenen Situationen zu wählen, wobei in jeder Situation die Trade-offs zwischen den eigenen und den Auszahlungen des anderen geändert werden kann. An Hand dieser sechs Wahlmöglichkeiten zeigen sie, wie man ein Maß an Prosozialität für jedes Subjekt, gemessen als Winkel und bestimmt durch die mittleren Zuweisungen an das Selbst und an die andere Person, zurückführen kann. Der Punktestand für soziale Werteorientierung steigt mit der Prosozialität der Wahlmöglichkeiten Werte unter 23 Grad signalisieren Wettbewerbsfähigkeit und Egoismus, höhere Werte eine eher prosoziale Disposition. Diese Aufgabe ist ein Standardinstrument zur Messung von sozialen Präferenzen, sowohl in der Sozialpsychologie als auch in der Ökonomie (vgl. Offerman et al., 1996, S. 823f.). Als zweites wird der Einfluss des Selbstwirkungsverständnisses mit Hilfe eines Fragebogens auf der Grundlage von Aquino et al. (2002, S. 1423f.) bewertet. In diesem Fragebogen wird die Testpersonen aufgefordert, die Attribute Fairness, Großzügigkeit und Freundlichkeit zu betrachten, und dann gebeten, bei sechs Aussagen, die von der Bedeutung der Attribute handeln, jeweils ihre Übereinstimmung oder Uneinigkeit mit ihrem Standpunkt aus auf einer Sechs-Punkte-Likert-Skala zu bewerten. Jede Behauptung kann auf einer Skala von 0 bis 5 bewertet werden, 0 entspricht "stark uneinig" und 5 "stark einig". Damit generiert man ein Maß für die Wichtigkeit des Selbstverständnisses, indem die Punkte jeder Aussage addiert werden. Grossman und van der Weele sammeln die Untersuchungsergebnisse von insgesamt 148 Teilnehmern, darunter 74 Diktatoren, die zusätzliche Einnahmen zwischen \$0,60 und \$4,00 an der Selbstwirkungsverständnis-Aufgabe verdienen. Der mitt-

lere Selbstwirkungsverständnis-Wert beträgt 35,9 Grad und der Median 34,1 Grad, beide in dem von Murphy, Ackermann und Handgraaf als prosozial beschriebenen Bereich (Murphy et al., 2011, S. 772f.). Beschränkt auf Diktatoren liegt der mittlere und der Median des Selbstwirkungsverständnis-Werts bei 35,9 und 34,5 Grad, was auf eine erfolgreiche Zufallsstreuung hindeutet (vgl. Tabelle 9). Der durchschnittliche und mediane Selbstimagegewichtungswert beträgt 22,9 bzw. 23 für die gesamte Stichprobe und 22,6 bzw. 23 für die Diktatoren. Um diese Zahlen in einen Zusammenhang zu rücken, folgen einige Beispiele. Ein Teilnehmer, der als Antwort auf alle sechs Aussagen über die Wichtigkeit des freundlichen, großzügigen und fairen Seins für sein Selbstwertgefühl "leicht zustimmen" auswählt, würde eine Selbstimagewertigkeit von 18 erhalten. Würde er allen sechs Aussagen "zustimmen", würde dies insgesamt eine 24 ergeben, bei „sehr zustimmen“ 30. Damit stimmt der durchschnittliche Teilnehmer den Aussagen zu. Die erste These macht Vorhersagen darüber, welche Typen die verschiedenen verfügbaren Optionen auswählen sollen. Die erste Vorhersage lautet, dass Diktatoren, die sich für Ignoranz entscheiden, als sozialer bewertet werden als solche, die sich für das Aufdecken des Spiels bei konträren Interessen entscheiden und dann egoistisch agieren und weniger sozial wahrgenommen werden als die Diktatoren, die sich für Transparenz und im Anschluss für eine faire Auszahlung entscheiden. Die zweite Vorhersage besagt, dass Diktatoren, die Transparenz wählen und sich dann bei konträren Interessen für eine eigennützig Auszahlung entscheiden, nachgesagt wird, dass sie sich weniger um ihre Selbstwirkung sorgen als Diktatoren, die entweder ignorant bleiben oder sich nach der Einforderung von Informationen prosozial bei konträren Interessen entscheiden. Der durchschnittliche Selbstwirkungsverständniswert für die 46 ignoranten Diktatoren liegt bei 34,0, nur knapp unter dem Durchschnittswert aller Diktatoren (vgl. Abbildung 6). Im Gegensatz dazu ist der Selbstwirkungsverständniswert der zehn Diktatoren, die sich erst für Aufdecken und dann bei konträren Interessen für B entschieden haben, signifikant höher als bei den ignoranten Diktatoren mit $p = 0,001$ beim einseitigen Mann-Whitney-U-Test (bei einem einseitigen parameterfreien Homogenitätstest). Obwohl die Tatsache, dass nur zwei Diktatoren, wissend, dass sie sich in der konträren Interessensituation befinden, egoistisch handeln, die Möglichkeit aus dem Vergleich starke Schlüsse zu ziehen stark limitiert, liegt der durchschnittliche Selbstwirkungsverständniswert dieser Diktatoren bei 28,36, was signifikant niedriger ist ($p = 0,04$, einseitiger Mann-Whitney-U-Test) als der der unwissenden Diktatoren. Damit unterstützt das Selbstwirkungsverständniswert-Maß die Sortierungsprognosen des Gleichgewichts. Nicht wahrnehmbar ist hingegen der Unterschied der Bewertung der ignoranten Diktatoren (22,54) und der Diktatoren, die sich bei konträren Interessen für B entscheiden (23,30). Im Gegensatz dazu ist die Selbstimage-Bewertung der beiden egoistischen Entscheider mit einem Wert von 12,50 deutlich niedriger als die der ignoranten Diktatoren ($p = 0,07$, einseitiger Mann-Whitney-U-Test). Auch diesmal schränkt der geringe Stichprobenum-

fang die Aussagekraft der Schlussfolgerungen ein, aber die großen Unterschiede legen nahe, dass Diktatoren in dieser Position weniger Bedenken hinsichtlich des Selbstbildes haben als diejenigen, die sich unwissend geben oder prosozial entscheiden. Schlussendlich unterstützen diese Ergebnisse die Überlegungen im Fall des Vergleichs über unterschiedliche Einsortierungsgruppen hinweg sowie die über Typen mit unterschiedlichen Verhaltensweisen (vgl. Grossman und Van der Weele, 2017, S. 191f.).

5.4.3. Rechtfertigung im Experiment

Um zu testen, ob Ignoranz als Entschuldigung dienen kann, wie in Abschnitt 5.3.3 angenommen, werden die Empfänger gebeten, den Charakter des Diktators, abhängig von der Wahl unterschiedlicher Strategien, zu bewerten. In neun Sitzungen des Hidden Information Experiments werden den Diktatoren sechs mögliche Handlungsabfolgen zur Verfügung gestellt: die Wahl von A oder B ohne Informationsempfangnis, Aufdecken und die Wahl von A oder B in der konträren Interessensituation und Aufdecken und die Wahl von A oder B bei gleichen Interessen. Für jede Variante beantworten die Empfänger die Frage "Wie sozial (im Gegensatz zu eigennützig) sehen sie Spieler X, wenn er oder sie die folgende Aktion auswählt?", indem sie einen Wert auf einer 5-Punkte-Skala von "sehr unsozial" bis "sehr sozial" auswählen. Die Empfänger vervollständigen diese Bewertungen, während die Diktatoren ihre Entscheidungen treffen. Die Skalenbewertung der Empfänger verläuft von 0 („sehr unsozial“) bis 4 („sehr sozial“). Abbildung 7 zeigt die Durchschnittsbewertungen der 72 Empfänger für Diktatoren, die sich für A oder B bei bewusster Ignoranz oder wissentlich unter der konträren Interessensituation entscheiden. Im Mittel schreiben Empfänger Diktatoren, die unter bewusster Ignoranz A wählen, eine soziale Bewertung von 1,69 zu. Deutlich höher als die, die A wissend wählen (1,10). Umgekehrt wird der Diktator, der informiert fair entscheidet (B), mit einem sozialen Ranking-Wert von 3,24 belohnt, während, wenn er sich unwissend dafür entscheidet, nur 2,35 Bewertungspunkte erhält. Der Mann-Whitney-Test zeigt, dass sich die Verteilung der Antworten in beiden Fällen auf einem 1%-Niveau signifikant unterscheidet. Damit kann Ignoranz in der Tat als Entschuldigung dienen (vgl. Grossman und Van der Weele, 2017, S. 194f.).

5.4.4. Bewusste versus unbewusste Ignoranz im Experiment

In Abschnitt 5.3.4 wurde gezeigt, dass die Versuchspersonen unter plausiblen Bedingungen stärker an Informationen über die Folgen ihres Handelns nach einer Handlung interessiert sein können als vor einer Handlung, auch wenn die Informationen danach nutzlos sind. Um diese Vorhersage zu testen, wird das Hidden Information Experiment dem Reveal Ex-Post Experiment, was sich im Wesentlichen dadurch unterscheidet, dass die Informationsentscheidung die zurückgestellte Auszahlungsentscheidung nicht mehr beeinflussen kann (vgl. modifizierte Spiel in Kapitel 5.3.4 bzw. Strategy Method in Kapitel 4.2), gegenüber gestellt. In Anlehnung an die 4. These wird davon ausgegangen, dass die Probanden im Reveal Ex-Post Experiment eher Informationen einfordern

als im Hidden Information Experiment. Darüber hinaus können die Ergebnisse des Reveal Ex-Post Experiments dafür verwendet werden zwei weitere Annahmen zu testen. Zum Ersten kann man erkennen, ob Individuen bereit sind den uninformierten Zustand zu überwinden und damit der Kritik entgegenzutreten. Ignoranz folgt hier aus der Standardauswahloption. Zum Zweiten spielen Informationen im Reveal Ex-Post Experiment keine instrumentelle Rolle mehr, so dass eine niedrige Ignoranzrate ein Beweis für Grossmans und van der Weeles Regel, dass die meisten indifferenten Probanden das Spiel aufdecken würden, wäre. Die Ergebnisse sind in Abbildung 8 dargestellt. Vergleicht man die beiden Experimente kann festgestellt werden, dass die Ignoranzrate im Reveal Ex-Post Experiment 26% beträgt und damit auf einem 1%-Niveau signifikant niedriger ist als die des Hidden Information Experiments mit 60%. Diese Tatsache stützt die Überlegung der beiden Forscher, dass Menschen generell lieber nach als vor der Auszahlungsentscheidung das Spiel aufdecken.

5.4.5. Bezahlung für Ignoranz im Experiment

In Abschnitt 5.3.5 wurde die Vorhersage abgeleitet, dass alle sozialen Typen, die einen niedrigeren Grad prosozialer Motivation als den Gleichgewichtswert besitzen, bereit sind, für Unwissenheit zu zahlen. In Folge dessen implizieren die komparativen Statistiken die steigende Ignoranz mit steigenden Informationskosten. Mit Einführung des Reveal Bonus Experiment werden diese beiden Vorhersagen getestet. In der Experimentumgebung können sich die Probanden \$0,10 dazu verdienen, wenn sie sich dazu entschließen, die Auszahlungsmatrix zu enthüllen. Mit dem Schriftzug „Spiel aufdecken + \$0,10“ werden sie gezielt auf diese Zusatzverdienstmöglichkeit aufmerksam gemacht. Durch die steigende Attraktivität von Informationen möchte man meinen, dass die Ignoranzrate sinkt. Als Analogie zur zweiten logischen Folgerung ist jedoch zu erwarten, dass die Ignoranzrate dennoch beträchtlich hoch ausfällt. Dies zeigt der rechte äußere Balken in Abbildung 8: die Ignoranzrate liegt bei 0,46. Obwohl die Ignoranzhäufigkeit niedriger ist als im Hidden Information Experiment, ist der Ergebnisunterschied bei Verwendung eines FET-Test ($p = 0.096$) nicht signifikant bei einem 10% Niveau. Trotz der Beweislage, dass Ignoranz minimal abnimmt, wenn sie kostspielig ist, ist in etwa die Hälfte der Probanden bereit dafür zu zahlen, weiterhin keine Informationen zu bekommen (vgl. Grossman und Van der Weele, 2017, S. 196).

5.5. Stellungnahme

Grossman und van der Weele können mit Hilfe ihrer Arbeit einen wesentlichen Einflussfaktor und Begründungsansatz für das Vermeiden von Informationen in unterschiedlichen Situationen finden. Es sind die Selbstwahrnehmung und die selbstkritischen Fragen darüber, wie andere einen wahrnehmen könnten, aufgrund des eigenen Wissenstands, des eigenen Handelns und zu Letzt des Verhaltens der Anderen (vgl. Grossman und Van der Weele, 2017, S. 176). Dass insbesondere zwischenmenschliche Interaktionen, wie auch der

Vergleich des eigenen Handelns mit dem der Anderen (vgl. 3. These), Einfluss auf das Ignoranzverhalten haben, kann Bénabou, der vorsätzliche Ignoranz in Gruppen untersucht, feststellen (vgl. 2012, S. 429f.). Damit wird der Begriff „bewusste Ignoranz“ in ein ganz neues Licht gestellt. Mit ihren Ergebnissen können Grossman und van der Weele zeigen, dass ein Gleichgewicht existiert, in dem es einen Grenztyp gibt, der sich für Ignoranz entscheidet, da er die Wechselwirkung zwischen seinem materiellen Vorteil und seinem Selbstbild in einer transparenten Situation fürchtet. Außerdem können sie ihre fünf verschiedenen Verhaltensmuster aus dem Gleichgewicht ableiten und diese in ihrem Experiment ganzheitlich beweisen und damit auch vorherige Forschungsergebnisse zum Teil bereichern, bestätigen und ergänzen (vgl. Tabelle 10). So hängt die Tatsache, dass Personen Informationen vermeiden, von unterschiedlichen Verhaltensstrukturen ab. Aus diesen leiten die Probanden situationsbedingt Entschuldigungen für Ignoranz und egoistische Verhaltenszüge her (vgl. Grossman und Van der Weele, 2017, S. 206f.). Dabei ist die Selbstwahrnehmung einer Person an sich einer der wichtigsten Einflussfaktoren für Verhaltensweisen in gesellschaftlichen Situationen und kann auch als Ursache für bewusste Ignoranz hervorgehoben werden. Personen haben oftmals regelrecht Angst davor, was andere von Ihnen denken könnten (vgl. Geiger und Swim, 2016, S. 79) und möchten den Erwartungshaltungen entsprechen (vgl. Dana et al., 2006, S. 199). Exemplarisch ist an dieser Stelle das immer wieder kritisch angesprochene Thema von Apothekern zu nennen, die gezielt Produktplatzierungen für die Pharmaindustrie betreiben. Apothekern wird viel Vertrauen entgegengebracht und sie besitzen ein hohes gesellschaftliches Ansehen. Geködert werden sie von der Pharmaindustrie, die sehr hohe Auszahlungen bietet (vgl. Pear, 2012, S. 1f.). Um nicht in einen moralischen Zwiespalt zu gelangen, werden gezielt Informationen über die Wirksamkeit von freiverkäuflichen Medikamenten vermieden (vgl. Kapitel 5.3.1). Je nach Vergütung für den Apotheker wird die Platzierung und Beratung des Kunden beeinflusst (vgl. Kapitel 5.3.5). Das eigene Interesse, sich über die Qualität von Alternativprodukten zu informieren, nimmt ab, was die Situation für aufklärungsmotivierte Firmen verschlechtert (vgl. Kapitel 5.3.2). Mit dem Verkauf des Produktes ändert sich für den Verkäufer die Haltung zu Informationen. Diese können seine Handlung nun nicht mehr beeinflussen (vgl. Kapitel 5.3.4). Und zu guter Letzt: auch ein Apotheker ist fehlbar. Diese Argumentation entschuldigt Fehlentscheidungen auf Grund von Unwissenheit (vgl. Kapitel 5.3.4). Nach jetziger Gewissheit darüber, dass Probanden zwischen ihrer Auszahlung und den möglichen Imagefolgen ihres Verhaltens abwägen, wäre es nun spannend in Erfahrung zu bringen, wie sie der Art beeinflusst werden, dass sie sich dennoch für mehr prosoziales Verhalten entscheiden. Trotz des Transfers auf eine reale Situation erweist sich das Paper von Grossman und van der Weele insbesondere durch ihr Modell als sehr abstrakt. Deshalb wird im nächsten Paper ein real-effort Experiment von Kajackaite analysiert. Diese Arbeit bringt einen Perspektivwechsel, da zuvor durch die Diktator-Spiel Methodik die Arbeitgeber bzw.

Handlungsträgerperspektive betrachtet wurde.

6. Die Auswirkungen von Unwissenheit auf Leistungen

In seiner Veröffentlichung „If I close my eyes, nobody will get hurt: the effect of ignorance on performance in a real-effort experiment“ behandelt Kajackaite die Thematik, ob Probanden abhängig von ihrer Entscheidung, sich über die Folgen ihres Handelns zu informieren, ihr Verhalten verändern. Ihr Verhalten wird auf Grund ihrer Leistung und der Kenntnis/Unkenntnis über den Spendenanteil, der sich abhängig von ihrem Verdienst vergrößert, an eine als negativ wahrgenommene Wohlfahrtsorganisation, gemessen. Wie erwartet, leisten Teilnehmer, die wissen, dass die NWO nichts erhält, mehr, als wenn sie wissen, dass diese anteilig Fördermittel erhält. Außerdem entscheidet sich ein Drittel, wenn es die Möglichkeit gibt zu erfahren, wie viel die NWO erhält, diese Information nicht zu erhalten. Generell leisten Personen, die sich aktiv für Ignoranz entscheiden, mehr als jene, die diese auferlegt bekommen (vgl. Kajackaite, 2015, S. 518).

6.1. Ignoranz unter realwirtschaftlichen Bedingungen

Der Blickwinkel aus Angestelltensicht ermöglicht es den Rahmen des Experiments nicht nur zu erweitern, sondern auch von der in den bisherigen Experimenten abstrakteren Vorgehensweisen ein „Ran-Zoomen“ an den industriellen Unternehmensalltag zu ermöglichen. Wie verhalten sich Angestellte, die über die Konsequenzen ihrer Arbeitsleistung nicht Bescheid wissen? Tendieren sie dazu sich zu informieren oder vermeiden sie, ebenso wie die Diktatoren, Informationen, die sich auf ihr Handeln auswirken könnten? Beeinflusst das Erfahren der negativen Folgen ihr Leistungsengagement? Mit der Beantwortung dieser Fragestellungen hebt sich die Arbeit von Kajackaite von bisherigen Forschungsergebnissen ab.

6.2. Vorstellung des Laborexperiments

In dem Experiment bekommen die Teilnehmer eine real-effort Aufgabenstellung. Sie müssen Buchstaben dekodieren, in dem sie für den entsprechenden Buchstaben die Zahl, die in derselben Zeile dahinter steht, in das Lösungskästchen eintragen. Erst nach dem Eintragen der richtigen Lösung kann der nächste Buchstabe dekodiert werden. Pro richtig identifizierten Buchstaben erhalten die Teilnehmer 5 ECU, wobei 100 ECU einem Euro entsprechen. Insgesamt gibt es vier verschiedene Experimentabwandlungen. Im Basisexperiment (BA) werden die Teilnehmer nach der Anzahl richtig dekodierter Buchstaben bezahlt. Das zweite Experiment (NRA) ist insofern anders, als die als negativ wahrgenommene Wohlfahrtsorganisation National Rifle Association bei jedem decodierten Buchstaben 7 ECU erhält. Die Probanden erhalten zusätzlich eine Beschreibung der NRA sowie die Resultate der Umfrageergebnisse des Campus zu der NRA, bei der 93% der Teilnehmer die Organisation als negativ empfanden. In dem dritten Experiment können sich die Personen aussuchen, ob sie die Auszahlung erfahren wollen oder nicht. Dabei besteht die Möglichkeit, dass die NRA entweder

keine Unterstützung in Abhängigkeit der Leistung oder pro gelösten Buchstaben 7 ECU erhält. Die Wahrscheinlichkeit für beide Ergebnisse ist 50%, da die Auszahlungsvariante vor Beginn eines Spiels per Münzwurf entschieden wird. Bei Teilnehmern, die sich für Transparenz entscheiden, werden die Fragezeichen dann durch die Auszahlungsmatrizen 0 ECU (NIO) oder 7 ECU (NI7) ersetzt. Die Teilnehmer, die ignorant bleiben, werden mit IG bezeichnet. Zusammengefasst werden die Ergebnisse der Spielteilnehmer, sowohl ignorant bleibende als auch nicht ignorante Spieler, unter der Kategorie NI&IG. Im letzten Experiment haben die Spieler den gleichen Kenntnisstand wie im vorherigen Experiment, jedoch haben sie nicht die Möglichkeit, die Auszahlungsvariante zu erfahren und verbleiben in Ungewissheit (UN). Die Opportunitätskosten werden in dem Spiel durch die Möglichkeit des Drückens eines Buttons dargestellt, mit dem das Spiel 20 Sekunden pausiert wird und der Teilnehmer 4 ECU erhält (vgl. [Kajackaite, 2015, S. 519f.](#)). Generell geht [Kajackaite](#) davon aus, dass die Probanden im Grundlagenexperiment eine positive Auszahlung generieren werden. Diese soll im zweiten Experiment aufgrund der als negativ wahrgenommenen Auszahlungen an die NRA geringer ausfallen. Außerdem wird, wie in anderen Experimenten ebenfalls bewiesen, erwartet, dass ein erheblicher Anteil sich im dritten Experiment für Ignoranz entscheiden wird (vgl. [Dana et al., 2007, S. 78f.](#); [Grossman, 2014, S. 2662f.](#)). Aufgrund des Einsortierungseffekts, der sich abhängig von einer endogenen oder exogenen Informationssituation ergibt, kann angenommen werden, dass Probanden, die sich bewusst für Informationen entscheiden, prosozialer agieren als Probanden, denen diese auferlegt wurde. Von Probanden, die in Unwissenheit bleiben müssen, wird jedoch erwartet, dass sie prosozialer auftreten als jene, die sich aktiv dafür entscheiden (vgl. [Lazar et al., 2012 S. 157.](#)). Außerdem wird davon ausgegangen, dass sich die Ergebnisse von BA und NIO nicht signifikant unterscheiden werden, da die finale Situation in beiden Experimenten dieselbe ist. Der Pausen-Button sollte nur dann interessant werden, wenn die Probanden Bescheid wissen, dass die NRA 7 ECU erhält und damit eine gute Alternative zu einem Arbeitsaufwand für die NRA erscheint. Insgesamt nehmen an den Experimenten 267 Probanden teil, jeder nur an einem sowie an einer der zehn Sessions. Um sicher zu gehen, dass die Probanden alles verstehen, bekommen sie vor Beginn der Experimente eine Einführung und dürfen Fragen stellen. Jedes Experiment startet mit einem 90 sekündigen Testlauf, der später als Fähigkeitsvergleich verwendet wird, bevor das eigentliche zehnminütige Experiment startet. Innerhalb der 90 Sekunden erhalten die Probanden jeweils 5 ECU für jeden dekodierten Buchstaben und können nicht den Pause-Button benutzen. Nach dem Experiment müssen die Probanden noch einen Fragebogen zu ihrer Person ausfüllen. Im Durchschnitt liegt die Auszahlung inklusive der 2,50€ Erscheinungsgebühr bei 12,81€ (vgl. [Kajackaite, 2015, S. 520f.](#)).

6.3. Auswertung der Laborergebnisse

Den Erwartungen entsprechend, erbringen die Probanden signifikant höhere Leistungen ($p = 0,01093$ unter Verwendung des parameterlosen zweiseitigen Fisher-Pitman Permutationstests, der im Folgenden bei Angabe des Signifikanzniveaus stets verwendet wird), wenn die NRA nichts erhält (BA: durchschnittlich 184,67 dekodierte Buchstaben), als wenn diese ebenfalls von den Leistungen profitiert (NRA: durchschnittlich 163,15 Buchstaben). Dass sich die Probanden unter den Bedingungen von NRA im Vergleich zu den BA Konditionen signifikant ($p = 0,02803$) unfairer behandelt fühlen/weniger wohl fühlen, zeigt ebenfalls eine Befragung nach Abschluss des Experiments, bei der die Probanden auf einer Likert Skala von 1 bis 7 angeben können wie fair sie die Konditionen empfinden (vgl. [Abbildung 11](#)). Bei der Wahl von Ignoranz entscheiden sich im dritten Experiment 28,25% der Probanden (36 von 127) dafür. Der Rest erfährt, wie viel die NRA erhält. Dabei ist der Leistungsunterschied zwischen der Gruppe, die erfährt, dass die NRA keine Auszahlung erhält, mit im Schnitt 190,33 dekodierten Buchstaben (NIO) signifikant höher ($p = 0,00234$) als die der Gruppe, die die Information erhält, dass die NRA 7 ECU pro richtige Lösung erhält (NI7), mit im Schnitt nur 150,69 Lösungen. Selbst der Anteil, der ignorant (IG) bleibt, erbringt im Schnitt mit 186,31 gelösten Buchstaben eine um 23,64% höhere Leistung als die NI7-Anteil, was bei einem Niveau von 0,01812 signifikant ist. Vergleicht man die Ergebnisse der Teilnehmer aus BA, NIO und IG miteinander, kann kein signifikanter Unterschied festgestellt werden. Dies zeigt, dass Ignoranz die Illusion von positiven Konsequenzen aufrecht erhält. Auch bei Probanden unter NRA und NI7 Bedingungen ist kein signifikanter Unterschied festzustellen. Die aktive Beschaffung bzw. die exogene Auflage der Information von negativen Konsequenzen hat keine Auswirkungen auf die Ergebnisse. Die Behauptung, dass ignorante Personen ein niedrigeres prosoziales Engagement zeigen als Personen, die sich nicht aktiv für Unwissenheit entscheiden (UN), kann bei marginaler Signifikanz ($p = 0,05757$) bestätigt werden. Sie sind in Unkenntnis der Konsequenzen des eigenen Handelns, was nicht unbedingt zu egoistischem Verhalten führen muss. Daraus kann man schließen, dass Personen mit einer geringen Neigung für prosoziales Handeln generell eine reserviertere Haltung gegenüber Angelegenheiten haben, sich auch nicht direkt ungerecht behandelt fühlen. Dem ist tatsächlich so, denn im Rahmen der Likert Skala geben Probanden der Gruppe IG einen bei einem Signifikanzniveau von $p = 0,05001$ niedrigeren Wert an. Das entspricht einer höheren Zufriedenheit als bei Probanden, die unwissend sind. Um zu schauen, ob die Fähigkeiten der Teilnehmer in den unterschiedlichen Experimenten gleich sind, werden die Ergebnisse des Testlaufs ausgewertet, die keinen signifikanten Unterschied ergeben. In Bezug auf das Nehmen von Pausen kann festgestellt werden, dass die Anzahl der Pausen bei NRA Bedingungen mit 3,25 Pausen signifikant höher ($p = 0,00005$) ist als der Durchschnitt im Grundlagenexperiment mit 0,10 Pausen. Diese Art der Arbeitsverweigerung kann als Alternative zu der Akzeptanz der Tatsache interpretiert

werden, dass die NRA an der eigenen Leistung ebenfalls profitiert (vgl. Abbildung 12). Dem entsprechend nehmen die Probanden in NI7 auch deutlich mehr Pausen ein als Agenten in NIO. So liegt die durchschnittliche Anzahl in NI7 und NIO bei 5,84 bzw. 0,07 Pausen ($p = 0,00011$). Die Agenten in IG nehmen durchschnittlich 1,03 Pausen ein, was deutlich weniger ist als im NI7-Zustand ($p = 0,01707$) (vgl. [Kajackaite, 2015, S. 221](#)).

6.4. Kritische Auseinandersetzung

Bei der Durchführung des real-effort Experiments, bei dem die Verhaltensweisen von Probanden auf negative Handlungskonsequenzen in einer realen Leistungssituation untersucht werden, ist auffällig, dass die Probanden, wenn sie die Möglichkeit haben Informationen zu vermeiden, dies weniger tun als in den Diktator-Spiel Varianten von DWK und Grossman (vgl. Tabelle 2 und Abbildung 4). Neu ist, dass sich in dem Fall die Ignoranz der als negativ empfundenen Konsequenzen nicht negativ auf das Leistungsverhalten auswirkt. Im Gegenteil die Probanden, die sich aktiv für Ignoranz entscheiden, dekodieren ähnlich viele Buchstaben wie Teilnehmer, die wissen, dass die NRA nichts erhält, zudem sind sie auch noch zufriedener. Dies liegt unter anderem daran, dass ihr prosoziales Verhalten und damit auch dessen Wirkung einen niedrigeren Stellenwert einnimmt (vgl. [Kajackaite, 2015, S. 523](#)). Umgekehrt verhält sich dies bei Spielern, die sich für Informationen entscheiden, dann aber erfahren, dass die NRA ebenfalls von ihrer Arbeit profitiert. Nicht nur, dass sie unzufriedener sind, sie verhalten sich zu dem noch ökonomisch ineffizient, in dem sie Pausen machen, obwohl ihnen in der Zwischenzeit schon allein ein dekodiertes Wort mehr Gewinn einbringen würde. Dies steht im Gegensatz zu dem in den Diktatorspielen festgestellten Verhaltensmuster der „Einsortierung“ (vgl. 2. These). Damit ist ein erster Hinweis darauf gegeben, dass Ignoranz bzw. das Erhalten von Informationen je nach Stellung des Spielers im System, ob arbeitend oder entscheidungstreffend, die Handlungsentscheidung auf ganz unterschiedliche Weise beeinflusst. [Kajackaite](#) zeigt, wie Ignoranz vor emotionalen Handlungen bewahren kann. Sie wird nicht allein von Spielern dazu genutzt, vor unangenehmen Konsequenzen die Augen schließen zu können, sondern auch vor sich selbst die Illusion von den gewünschten Handlungsauswirkungen aufrecht erhalten zu können. Da wie eben beschrieben, Ignoranz in der Experimentssituation, um unter anderem eine rationale Haltung gegenüber der Aufgabenstellung wahren zu können, sehr attraktiv ist, stellt sich einem die Frage, warum hier Ignoranz nur von knapp mehr als einem Viertel gewählt wurde. Wünschenswert wäre gewesen, ebenfalls andere Wahlmethoden von Ignoranz wie z.B. der Strategy Method einzuführen, eine Möglichkeit um herauszufinden, ob es situationsbedingte Verhaltensunterschiede auch in real-effort Experimenten gibt. Außerdem könnten weitere Unterschiede zu Experimenten, in denen keine Leistung erbracht werden muss, erforscht werden.

7. Schlussbetrachtung und Fazit

Ziel der Arbeit war es auf zu decken, ob Probanden gezielt Informationen vermeiden, die eine nachfolgende Handlungsentscheidung beeinflussen könnten. Diese Tatsache konnte in allen vier in der Arbeit analysierten Publikationen bestätigt werden. Auch situative Einflüsse konnten konkretisiert werden. So werden Informationen im intransparenten Handlungsrahmen im Schnitt zu 25% mehr vermieden (vgl. Tabelle 7) als bei vollständigem Kenntnisstand. Dabei kann die Standardauswahl mit kleinen methodischen Änderungen der Informationswählerhebung eine entscheidende Rolle spielen (vgl. [Grossman, 2014, S. 2660](#)). So bevorzugen Probanden intransparente Situationen, in denen die Informationsentscheidung passiv wählbar ist. Begründet werden kann die Entscheidung des Spielers durch die individuelle Nutzenfunktion, in der er Imagevorteile und materiellen Nutzen gegeneinander, in Abhängigkeit von seinem Grad prosozialer Motivation, abwägt. Dennoch hat die Vermeidung von Informationen unterschiedliche Konsequenzen, was sich durch die Anwendung unterschiedlicher Spielvarianten zeigt. Dadurch ergeben sich verschiedenartige Lösungsansätze. Die Diktatorspielvariante zeigt einen Handlungsträger, der mit seiner Entscheidung die Situation seiner Mitspieler, ohne deren Einwilligung, aktiv gestalten kann. Diese Verhältnisse sind in der Konstellation Arbeitgeber zu Arbeitnehmer anzutreffen. Die Gewerkschaften nehmen seit Ende des 19. Jahrhunderts Einfluss auf dieses Spannungsverhältnis. Dass diese jedoch das bewusste Vermeiden von Informationen nicht ausreichend unterbinden können, zeigt die VW-Abgas-Affäre, bei der vor zwei Jahren bekannt wurde, dass manipulierte Software verwendet wurde. Ein großer Schaden für das Unternehmen. DWK's Baseline Experiment zeigt, dass 74% der Probanden fair handeln (vgl. Tabelle 1), wenn sie mit vollen Informationen ausgestattet werden. Damit wäre die Einstellung eines vom Staat bezahlten unabhängigen Mitarbeiters eine Möglichkeit, Führungskräfte mit Informationen zu konfrontieren, da insbesondere soziale Präferenzen eine größere Wirkung haben, wenn Informationen über die Folgen der eigenen Entscheidung nicht passiv vermieden werden können (vgl. [Grossman, 2014, S. 2660](#)). Der freie Mitarbeiter würde zu dem, losgelöst von der Unternehmensethik, die das Verhalten eines Mitarbeiters maßgeblich beeinflusst, handeln (vgl. [Pierce und Snyder, 2008, S. 1900](#)). Dieser eben präsentierte Lösungsansatz hat nur eine kleine Kongruenz mit dem weit verbreiteten Begriff des Whistleblowings, da es nicht die Intention ist, Skandale aufzudecken und dem Unternehmen zu schaden, sondern unternehmensintern Probleme in Angriff zu nehmen. Geht es hingegen um die Mitarbeitermotivation, so veranschaulicht das real-effort Experiment von [Kajackaite \(2015\)](#), dass es den Mitarbeitern selbst überlassen sein sollte, ob sie Konsequenzen ihrer Unternehmensleistung erfahren wollen. Wenn das Wissen dazu führen sollte, dass sie sich aufgrund der von ihnen als negativ empfundenen Konsequenzen weniger einbringen, wirkt sich dies auch auf das Unternehmen negativ aus. Genauso verhält es sich hinsichtlich des sozialen Engagements. Auch dort ist es den Menschen über-

lassen, wie stark sie sich selbst einbringen (vgl. Niehaus, 2013, S. 3). Selbst wenn sie nicht erfahren wollen, welchen Wirkungsgrad ihre Spende erzielt, stellt sich die Frage: Sollte man ihnen ihre prosoziale Motivation nehmen, in dem man Ihnen Informationen „aufdrängt“? An diesen Gedanken knüpft auch die aktuelle Forschung von Kandul und Ritov an. Sie zeigen, dass bewusste Ignoranz auch prosozial genutzt werden kann. In einem Dualen-Selbst-Paradigma kann die egoistische Hälfte einer Person auf der einen Seite bewusste Ignoranz gegen die prosoziale Hälfte verwenden, auf der anderen Seite kann die prosoziale Hälfte mit den gleichen Mitteln die egoistische Hälfte vor Versuchungen bewahren (vgl. Kandul und Ritov, 2017, S.56). Die gezielte Nutzung von bewusster Ignoranz um prosoziales Verhalten hervorzurufen, unternehmensbezogener Diktatorspielvarianten sowie die Untersuchung der Wirkung von einer Aufsichtsperson im Unternehmen wären in weiterführenden Forschungen wünschenswert.

Die Tatsache, dass Ignoranz an der Wahrheit nichts ändern kann, erkannte bereits Winston Churchill: „The truth is incontrovertible. Malice may attack it, ignorance may deride it, but in the end, there it is.“ (Miller, 2009, S.243).

Literatur

- Admati, A. R. It takes a village to maintain a dangerous financial system, 2017. URL <https://ssrn.com/abstract=2787177>. Working Paper. Stable SSRN:.
- Aghion, P. und Tirole, J. Formal and real authority in organizations. *Journal of political economy*, 105(1):1–29, 1997.
- Aquino, K., Reed, I., et al. The self-importance of moral identity. *Journal of personality and social psychology*, 83(6):1423, 2002.
- Ariely, D., Bracha, A., und Meier, S. Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review*, 99(1):544–555, 2009.
- Bandura, A., Barbaranelli, C., Caprara, G. V., und Pastorelli, C. Mechanisms of moral disengagement in the exercise of moral agency. *Journal of personality and social psychology*, 71(2):364–374, 1996.
- Barfort, S., Harmon, N. A., Hjorth, F. G., und Olsen, A. L. Dishonesty and selection into public service in denmark: Who runs the world's least corrupt public sector? *Discussion Paper*, pages 1–57, 2015.
- Bateman, T. S. und O'Connor, K. Felt responsibility and climate engagement: Distinguishing adaptation from mitigation. *Global environmental change*, 41:206–215, 2016.
- Bénabou, R. und Tirole, J. Incentives and prosocial behavior. *American economic review*, 96(5):1652–1678, 2006.
- Broberg, T., Ellingsen, T., und Johannesson, M. Is generosity involuntary? *Economics Letters*, 94(1):32–37, 2007.
- Cartwright, E. und Menezes, M. L. Cheating to win: Dishonesty and the intensity of competition. *Economics Letters*, 122(1):55–58, 2014.
- Casari, M. und Cason, T. N. The strategy method lowers measured trustworthiness behavior. *Economics Letters*, 103(3):157–159, 2009.
- Chakravarty, S., Harrison, G. W., Haruvy, E. E., und Rutström, E. E. Are you risk averse over other people's money? *Southern Economic Journal*, 77(4):901–913, 2011.
- Charness, G. und Jackson, M. O. The role of responsibility in strategic risk-taking. *Journal of Economic Behavior & Organization*, 69(3):241–247, 2009.
- Claassen, U. Die wahren Ursachen der VW-Krise, 2017. URL <http://www.manager-magazin.de/unternehmen/autoindustrie/volkswagen-die-wahren-ursachen-der-krise-a-1159434-3.html>. Stand: 29.12.2017.
- Conrads, J., Irlenbusch, B., Rilke, R. M., und Walkowitz, G. Lying and team incentives. *Journal of Economic Psychology*, 34:1–7, 2013.
- Crémer, J. Arm's length relationships. *The Quarterly Journal of Economics*, 110(2):275–295, 1995.
- Dana, J., Cain, D. M., und Dawes, R. M. What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100(2):193–201, 2006.
- Dana, J., Weber, R. A., und Kuang, J. X. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80, 2007.
- Dawson, E., Savitsky, K., und Dunning, D. "don't tell me, i don't want to know": Understanding people's reluctance to obtain medical diagnostic information 1. *Journal of Applied Social Psychology*, 36(3):751–768, 2006.
- Drewes, A. Stellungnahme zu schmutzige Schokolade II, 2012. URL <https://www.nestle.de/medien/medienstatements/stellungnahme-zu-schmutzige-schokolade-ii>. Stand: 05.01.2018.
- Dutton, D. G., Boyanowsky, E. O., und Bond, M. H. Extreme mass homicide: From military massacre to genocide. *Aggression and Violent Behavior*, 10(4):437–473, 2005.
- Ehrich, K. R. und Irwin, J. R. Willful ignorance in the request for product attribute information. *Journal of Marketing Research*, 42(3):266–277, 2005.
- Exley, C. L. Excusing selfishness in charitable giving: The role of risk. *The Review of Economic Studies*, 83(2):587–628, 2015.
- Exley, C. L. und Kessler, J. B. Why information fails to encourage costly behavior. Technical report, Working paper, 2017.
- Feiler, L. Testing models of information avoidance with binary choice dictator games. *Journal of Economic Psychology*, 45:253–267, 2014.
- Firestein, S. *Ignorance: How it drives science*. New York: Oxford University Press Inc., 2012.
- Fischbacher, U., Gächter, S., und Quercia, S. The behavioral validity of the strategy method in public good experiments. *Journal of Economic Psychology*, 33(4):897–913, 2012.
- Galperin, B. L., Bennett, R. J., und Aquino, K. Status differentiation and the protean self: A social-cognitive model of unethical behavior in organizations. *Journal of Business Ethics*, 98(3):407–424, 2011.
- Geiger, N. und Swim, J. K. Climate of silence: Pluralistic ignorance as a barrier to climate change discussion. *Journal of Environmental Psychology*, 47:79–90, 2016.
- Gneezy, A., Gneezy, U., Riener, G., und Nelson, L. D. Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences*, 109(19):7236–7240, 2012.
- Golman, R., Hagmann, D., und Loewenstein, G. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, 2017.
- Grossman, Z. Self-image and strategic ignorance in moral dilemmas. *Working Paper*, 2013.
- Grossman, Z. Strategic ignorance and the robustness of social preferences. *Management Science*, 60(11):2659–2665, 2014.
- Grossman, Z. Self-signaling and social-signaling in giving. *Journal of Economic Behavior & Organization*, 117:26–39, 2015.
- Grossman, Z. und Van der Weele, J. J. Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, 15(1):173–217, 2017.
- Hamman, J. R., Loewenstein, G., und Weber, R. A. Self-interest through delegation: An additional rationale for the principal-agent relationship. *American Economic Review*, 100(4):1826–1846, 2010.
- Hobson, K. und Niemeier, S. "what sceptics believe": The effects of information and deliberation on climate change scepticism. *Public understanding of science*, 22(4):396–412, 2013.
- Johnson, E. J. und Goldstein, D. Do defaults save lives? *Science*, 302(5649):1338–1339, 2003.
- Kajackaite, A. If i close my eyes, nobody will get hurt: The effect of ignorance on performance in a real-effort experiment. *Journal of Economic Behavior & Organization*, 116:518–524, 2015.
- Kajackaite, A. und Gneezy, U. Incentives and cheating. *Games and Economic Behavior*, 102:433–444, 2017.
- Kandul, S. und Ritov, I. Close your eyes and be nice: Deliberate ignorance behind pro-social choices. *Economics Letters*, 153:54–56, 2017.
- Krupka, E. L. und Weber, R. A. Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association*, 11(3):495–524, 2013.
- Larsen, S. V., Kørnøv, L., und Driscoll, P. Avoiding climate change uncertainties in strategic environmental assessment. *Environmental Impact Assessment Review*, 43:144–150, 2013.
- Larson, T. und Capra, C. M. Exploiting moral wiggle room: Illusory preference for fairness? a comment. *Judgment and Decision Making*, 4(6):467–474, 2009.
- Larson, T. C. und Noussair, C. The use of strategic ignorance in dictator games when payoffs are not-transparent. *Department of Economics Master's Thesis, Emory University*, 2005.
- Lazear, E. P., Malmendier, U., und Weber, R. A. Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics*, 4(1):136–163, 2012.
- Madrian, B. C. und Shea, D. F. The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics*, 116(4):1149–1187, 2001.
- Mazar, N., Amir, O., und Ariely, D. The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research*, 45(6):633–644, 2008.
- Miller, J. Emotion in politics: Blue and red brains. *The Review of Communication*, 9(3):243–246, 2009.
- Mistrati, M. Das Erste, 2010. URL <http://www.daserste.de/information/reportage-dokumentation/dokus/videos/schmutzige-schokolade-100.html>. Stand: 05.01.2018.
- Mitzkewitz, M. und Nagel, R. Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22(2):171–198, 1993.
- Murnighan, J. K. und Wang, L. The social world as an experimental game. *Organizational Behavior and Human Decision Processes*, 136:80–94, 2016.
- Murphy, R. O., Ackermann, K. A., und Handgraaf, M. Measuring social value orientation. *Judgment and Decision making*, 6(8):771–781, 2011.
- Niehaus, P. A Theory of Good Intentions, 2013. URL http://cgeg.sipa.columbia.edu/sites/default/files/cgeg/S13_Niehaus_0.pdf. Stand: 20.01.2018.
- Norgaard, K. M. "we don't really want to know" environmental justice and

- socially organized denial of global warming in norway. *Organization & Environment*, 19(3):347–370, 2006.
- Nyborg, K. I don't want to hear about it: Rational ignorance among duty-oriented consumers. *Journal of Economic Behavior & Organization*, 79(3): 263–274, 2011.
- Offerman, T., Sonnemans, J., und Schram, A. Value orientations, expectations and voluntary contributions in public goods. *The economic journal*, 106(437):817–845, 1996.
- Paharia, N., Vohs, K. D., und Deshpandé, R. Sweatshop labor is wrong unless the shoes are cute: Cognition can both help and hurt moral motivated reasoning. *Organizational Behavior and Human Decision Processes*, 121(1):81–88, 2013.
- Pear, R. U.S. to Force Drug Firms to Report Money Paid to Doctors, 2012. URL <http://www.nytimes.com/2012/01/17/health/policy/us-to-tell-drug-makers-to-disclose-payments-to-doctors.html>. Stand: 07.01.2012.
- Pierce, L. und Snyder, J. Ethical spillovers in firms: Evidence from vehicle emissions testing. *Management Science*, 54(11):1891–1903, 2008.
- Sahlane, A. Argumentation and fallacy in the justification of the 2003 war on iraq. *Argumentation*, 26(4):459–488, 2012.
- Sweeny, K., Melnyk, D., Miller, W., und Shepperd, J. A. Information avoidance: Who, what, when, and why. *Review of general psychology*, 14(4): 340–353, 2010.
- Thunström, L., van't Veld, K., Shogren, J. F., und Nordström, J. On strategic ignorance of environmental harm and social norms. *Revue d'économie politique*, 124(2):195–214, 2014.
- Thunström, L., Nordström, J., Shogren, J. F., Ehmke, M., und van't Veld, K. Strategic self-ignorance. *Journal of Risk and Uncertainty*, 52(2):117–136, 2016.
- Van der Weele, J. Inconvenient truths: Determinants of strategic ignorance in moral dilemmas."working paper. 2014.



The Effect of ECB's Corporate Sector Purchase Programme on CDS Premia - An Empirical Analysis

Silie Homayon Nawabi

Goethe Universität Frankfurt

Abstract

In response to the intensification of economic crises in the euro area, the European Central Bank (ECB), along with other central banks, has conducted both conventional and unconventional monetary policy. The most recent unconventional measure has been outright asset purchases under the corporate sector purchase programme (CSPP) targeting euro-denominated investment-grade bonds issued by non-financial corporations in the euro area. Using a Difference-in-Differences (DID) approach on a sample of euro-zone data I find that the CSPP initiative has consistently contained credit risk. In contrast, spillover effects to firms not subject to the CSPP policy are limited.

Keywords: quantitative easing; unconventional monetary policy; asset purchase program; credit default swaps; corporate sector purchase program

1. Introduction

'Likewise, the credit easing components of our expanded asset purchase programme (APP), namely the asset-backed securities (ABSPP), covered bond (CBPP3) and corporate sector (CSPP) purchase programmes, further boost the pass-through of our monetary policy by directly lowering the financing costs for crucial actors in our economy. [...] the CSPP directly lowers the cost and improves the availability of market-based funding for non-financial corporations.' - Mario Draghi, ECB President, Brussels, 26 September 2016

After the failure of Lehman Brothers in September 2008, confidence in the world economy collapsed and international financial markets became severely disrupted. Soon the European sovereign debt crisis followed, posing further challenges for the euro area. In response, the ECB not only implemented a drastic cut in its official interest rates, but also introduced a package of non-standard monetary policy measures. These measures were motivated by the need to ensure the continued effectiveness of the transmission of the monetary policy stance to the real economy and ultimately to price developments (Giannone et al., 2011). Yet in Europe, financial markets remained dysfunctional with credit condi-

tions tightening markedly, and the risk of depressed inflation rates. As a complement to existing unconventional measures the CSPP has been launched, with a first formal announcement in March 2016. In effect, the ECB has expanded its quantitative easing (QE) programme to include the purchase of non-financial corporate bonds. The CSPP commenced officially on June 8th, 2016, with the objective to provide further monetary policy accommodation and contribute to a return of inflation rates to levels below, but close to two percent (European Central Bank, 2016). On average seven billion euro corporate bonds are bought each month. By the end of November 2017, CSPP purchases had reached a total of € 128 billion.

According to the ECB, this QE programme accounts for a large decline in funding costs for both financial and non-financial corporations in the eurozone. At the same time, the programme contributes to the bypassing of financial intermediaries by an increased availability of credit. Thus, by means of the CSPP, the ECB has been able to lift credit constraints notably, in an environment where the financial system has been subject to considerable stress (European Central Bank, 2016). Empirical research, placing a special focus on the corporate bond market, further confirms the effective transmission of the monetary policy to the real economy. While the response of cash market instruments to unconventional monetary policy measures has been investigated thoroughly in the

literature, especially for those measures longer in place, evidence on derivative markets is scarce. This paper contributes to the literature by documenting the effect of the CSPP on credit derivatives.

Indeed, it is inappropriate to focus only on the cash market when assessing the CSPP impact. As emphasized by Krishnamurthy and Vissing-Jorgensen (2011) and Eser and Schwaab (2016), QE programmes work through various channels. In this paper I will particularly address the default risk channel and the portfolio rebalancing channel. First, given the CSPP succeeds in stimulating the economy by lowering borrowing costs for corporates, one should observe a reduction in expected defaults and, as a result, a decline in corporate credit risk. Moreover, as the economic recovers, standard asset pricing models imply - beyond the compensation for expected defaults - a reduction in the average price for assuming exposure to corporate credit risk. In fact, investors' risk aversion is expected to decline, implying a lower default risk perception, and ultimately a lower default risk premium (Gilchrist and Zakrajšek, 2013). My second conjecture is that the CSPP policy may contribute endogenously through spillover effects, in line with the theory of the portfolio rebalancing channel (Altavilla et al., 2015). The sizeable compression of funding costs induced by the CSPP should be reflected in substantially lower costs of default insurance, especially for riskier credits. More specifically, this line of argument suggests that, while credit risk has reduced overall, the policy impact is not restricted to CSPP-eligible assets but extends beyond the eligibility criteria.

The objective of this paper is to quantify these impacts undertaken within the CSPP framework on market-based measures of corporate credit risk; in particular on credit default swap (CDS) spreads. In essence, I assess whether CDS show price reactions consistent with the intentions behind the monetary policy strategy of the ECB. The market of single name CDS is of particular interest because, by their nature, these innovative instruments equip researchers with a near-ideal way of directly measuring credit risk (Longstaff et al., 2005, Norden and Weber, 2009). In general, CDS¹ are bilateral contracts that provide protection against the risk of a credit event associated with a particular company or country. Hence, they serve as a vehicle through which market participants are able to isolate and transfer credit risk (Fontana and Scheicher, 2016).

Using a panel of eurozone CDS data, I seek to empirically investigate the effect of the CSPP event on firms' CDS spreads within a DID framework. The empirical strategy identifies the distinct CSPP purchase dates as the most important piece of information to causally link the CSPP with the outcome of interest. In fact, by exploiting the phased implementation of

the CSPP policy, I am able to address the endogeneity concern of non-random assignment of CSPP-eligible bonds. In other words, the distinct purchase dates allow a comparison between the subsample of firms transferred primarily to the CSPP portfolio (treatment group) and firms transferred later (control group). In this vein, the full sample is restricted to CSPP firms only. The within CSPP-sample analysis then mitigates any concerns related to heterogeneity within the treatment group in response to the CSPP, and ultimately any concerns related to omitted variables. Apart from that, the DID estimation is undertaken on a set of firms within the same industry.

Consistent with the initial assumption, I find that the CSPP programme has contained credit risk across European non-financial corporates. The results indicate that credit market reactions to the CSPP event - measured by means of CDS prices - imply negative CDS rates throughout. The most pronounced impact in lowering credit risk can be observed for the sector of Basic Materials, accounting for a CDS spread decrease of approximately 8 percent. In contrast, the empirical support for the second conjecture is limited. Spillover effects to firms not yet subject to the CSPP are heterogeneous within and across industries, and if anything, rather bond specific. For example, within the Industrial sector, I observe considerable spillover effects of around 6 percentage points for a given reference firm. However, a precedent bond purchase referring to the same firm does not prompt any spillovers. Hence, the ECB's commitment to continue the CSPP is indeed helping to lift credit constraints overall, but according to my estimation the initiative seems to not have stronger effects on firms unaffected by the CSPP, as suggested by previous work.

The rest of this paper is organized as follows. Section 2 discloses the ECB's monetary policy strategy as a response to the global financial and the European sovereign debt crisis, with a special focus on the CSPP initiative. Section 3 provides a brief overview over corporate CDS and highlights their importance as a corporate credit risk measure. Section 4 introduces two hypotheses as well as describes the underlying empirical strategy and the sample data. In Section 5 the main empirical findings of the CSPP impact on corporate CDS are presented and discussed. The paper closes with a discussion in Section 6.

2. A New Wave of Unconventional Monetary Policy

The Governing Council of the ECB assesses economic and monetary developments and takes monetary policy decisions every six weeks (European Central Bank Website, 2018b). The primary objective of its monetary policy stance is to maintain price stability within the Eurozone. In particular, price stability is defined as a year-on-year increase in the Harmonised Index of Consumer Prices of below two percent (European Central Bank, 2011). Recent economic shocks have posed significant challenges for the euro area. In response, the ECB has not only cut its official interest rates

¹Unlike multiname CDS, the underlying of single-name CDS refers to a single firm or entity. Multi-name CDS such as basket CDS or CDS indices, on the other hand, are written on a set of firms. Analysing multi-name CDS is beyond the scope of the present study. As my focus is on singlename CDS, in the following, I will use the terms single-name CDS and CDS interchangeably.

significantly, but also has adopted a series of unconventional monetary policy measures such as the CSPP.

2.1. ECB's Unconventional Monetary Policy Measures

In general, conventional monetary policy operates by steering nominal short-term interest rates at which commercial banks can borrow funds from and deposit funds at the central bank (Joyce et al., 2012). The underlying economic rationale is to lower these key interest rates during economic downturns and to increase them during economic upturns.² In this manner, the central bank effectively manages the liquidity conditions in money markets. For decades the ECB has successfully relied on this standard interest rate channel to fulfil its price stability mandate over the medium term.

However, such measures are no longer sensible when interest rates are already close to the zero bound. The zero lower bound describes the notion that interest rates cannot be below zero percent. If so, agents in the economy would hold zero interest cash instead (Keynes, 1936). Hence, as soon as interest rates are close to the effective lower bound, conventional interest rate targeting will cease to be effective and central bank authorities will have to opt for unconventional measures to stabilize price levels in particular and the economy in general (Eggertsson and Woodford, 2003, Bernanke and Reinhart, 2004, Hamilton and Wu, 2012, Woodford, 2012).

Indeed, in the wake of the global financial crisis - soon followed by the sovereign debt crisis in several euro area countries - interest rates quickly approached the effective lower bound. Nonetheless, given the scale of losses incurred in the aftermath of the crisis the financial system remained dysfunctional. In fact, shortly after the collapse of Lehman Brothers, interbank market liquidity virtually dried up. Banks abandoned making loans and asset prices dropped dramatically leaving the financial system as a whole exposed to the risk of a liquidity trap, in which each economic agent is keen to hoard liquidity (Beirne et al., 2011, Joyce et al., 2012). The spread between the three-month risky interbank rate (EURIBOR) and the overnight interest rate (EONIA) - a key factor to evaluate the health of the European interbank market - rose to a new all-time high of 156 basis points in October 13, 2008 (Bini Smaghi, 2009). This ongoing tension within financial markets left the monetary policy transmission process severely impaired. Thus, the ECB was eager to provide additional monetary stimulus to the economy beyond the standard interest rate channel (Joyce et al., 2012, Fawley and Neely, 2013).

Broadly speaking, unconventional measures are defined as those policies that directly cope with funding needs of banks, households and non-financial companies. Financial support by the central bank authority can be provided in the form of central bank liquidity, loans, fixed-income securities

or equity. Principally, as the cost of external finance is traded at a premium, the set of unconventional measures can be regarded as an attempt to reduce specific risk premia. Particularly, the central bank may reduce term spreads between short and long-term rates and/or credit spreads between risk-free assets and risky assets, eventually influencing wealth, cost of borrowing, spending and income (Joyce and Tong, 2012, Mertens, 2017).

In general, central bank authorities can select from a wide range of unconventional measures which are not necessarily mutually exclusive. Typically, they also serve as a complement to standard interest rate decisions rather than to substitute for them (Giannone et al., 2011). The final choice depends on institutional features, the structure of the financial system, the degree of disruption within markets and most importantly the intermediate objectives. A stylized representation of potential measures is shown in Figure 1. As demonstrated, unconventional monetary policy can be allocated to two broad categories: forward guidance or balance sheet measures. Forward guidance represents the central bank's commitment to the public to maintain its accommodative monetary policy over an extended period, such as to keep short-term interest rates low for a significant period of time.³ In fact, the speech by the ECB's president Mario Draghi on July 26, 2012, in which he stated that 'the ECB is ready to do whatever it takes to preserve the Euro', may be interpreted as a forward guidance tool (Draghi, 2012).

In stark contrast, balance sheet measures affect explicitly the size or the composition of the central bank's balance sheet, officially known as quantitative easing or credit easing. The subset can be further differentiated into direct and indirect measures. With direct measures in place, the central bank engages in direct acquisition of assets, until maturity or resale, and thus assumes the associated risks⁴ on its balance sheet. In standard literature outright asset purchase programmes such as the CSPP are identified as a direct QE policy, see for example Draghi (2016), Abidi et al. (2017) and Arce et al. (2017). Through the indirect approach the central bank lends to other banks at longer maturities in exchange for collateral, including assets whose markets are temporarily impaired. By this means the central bank does not assume any risks on its balance sheet (Woodford, 2012). Usually, indirect measures lead to a relatively small or even no increase of the central bank's balance sheet, whereas asset purchase programmes within the QE framework are undertaken in large and highly liquid market segments, that in turn lead to a substantial expansion of the central bank's balance sheet (Mertens, 2017). A deeper investigation into each unconventional measure conducted by the ECB is beyond the scope of this paper. For a detailed overview, see the contributions by the former ECB executive board member Lorenzo Bini Smaghi (Bini Smaghi, 2009).

²In particular, the ECB sets the target overnight interest rate in the interbank money market, in this manner, signaling the desired policy rate. Hereof, the prominent Taylor rules provide guidelines regarding the level of short-run benchmark rates (Taylor, 1993).

³In principle, this communicative instrument can be conditional or unconditional. For more details see Bernanke and Reinhart (2004).

⁴Potential risk sources to the balance sheet may materialize through interest rate risk, market risk, sovereign risk or credit risk.




Figure 1: Unconventional Monetary Policy; Source: [Bernanke and Reinhart \(2004\)](#), [Bini Smaghi \(2009\)](#), [Draghi \(2016\)](#).

Unconventional monetary policy is based on the idea that the central bank can stimulate the economy, even when short-term interest rates are at or close to zero. The figure shows the different measures that central banks may adopt. Broadly, these measures can be differentiated between forward guidance and balance sheet measures. All policies related to quantitative easing and credit easing are a subset of balance sheet measures.

Given that the European Economic and Monetary Union (EMU) is a bank-based economy, the first set of unconventional measures adopted by the ECB - officially known as the 'Enhanced Credit Support' - was directed at banks ([Fawley and Neely, 2013](#), [Mertens, 2017](#)). In this context, the former ECB President Jean-Claude Trichet stressed that the ECB would neither involve in direct credit easing nor in direct quantitative easing, as implemented by other major central banks at that time, but focus on endogenous measures.⁵ The policy comprised five building blocks: fixed-rate full-allotment (FRFA), expansion of eligible collateral, longer-term liquidity provision, liquidity provision in foreign currencies and financial market support through purchases of covered bonds (CBPP1).⁶ At some later stage, the ECB also decided to offer accommodative refinancing facilities for banks over longer periods of time by means of its longer-term re-

⁵Japan has become known for its QE policy with the expansion of the monetary base through outright purchases of government bonds from the banking sector. By contrast, the Bank of England has bought British government bonds from the non-bank private sector. The US Federal Reserve was initially rather engaged in credit easing by providing direct lending facilities to market participants but over the course of time it has resorted to QE initiatives as well, buying securities from government agencies ([Trichet, 2009](#), [Joyce et al., 2012](#)).

⁶For instance, the CBPP1 as an indirect measure was initiated to alleviate the potential risk of a bank run given the maturity mismatch banks are exposed to when granting long-term loans, financed by short-term deposits ([Joyce et al., 2012](#), [Fawley and Neely, 2013](#)).

financing operations (LTRO) programme. Overall, this indirect approach was intended to primarily alleviate growing tensions in interbank money markets, to expand bank-lending operations and ultimately to ensure the transmission of the ECB's policy stance to the real economy ([Trichet, 2009](#), [Joyce et al., 2012](#), [Fawley and Neely, 2013](#), [Szczerbowicz et al., 2015](#), [Eser and Schwaab, 2016](#)).

Yet, in spite of these efforts, the European debt crisis deepened further as banks were heavily exposed to risky sovereign debt issued by periphery Eurozone countries ([Szczerbowicz et al., 2015](#)). At its height in 2012, the ECB eventually decided to extend its unconventional policy toolkit and resort to direct measures. The ECB was able to calm financial markets by announcing conditional support by means of a sovereign state bailout programme, namely Outright Monetary Transactions (OMT). In addition, purchases in sovereign debt markets within the Securities Markets Programme (SMP) framework had been announced. Nonetheless, the SMP was not deliberately designed as a QE-type programme ([Eser and Schwaab, 2016](#), [Schlepper et al., 2017](#)).

The trend-breaking effect in ECB's monetary policy stance was actually induced in January 2015, with the introduction of the expanded asset purchase programme (APP). Eventually, the ECB joined several other central banks in implementing QE with outright purchases of EMU government bonds on an unprecedented scale. Since then the programme has been

modified and amended several times. For example, in October 2017, the public sector purchase programme (PSPP) - as one pillar of the APP - was reduced from a monthly purchase pace of € 60 billion to € 30 billion (Andrade et al., 2016, European Central Bank, 2017b). While the APP policy was initially launched to restore the smooth functioning of financial markets, concerns soon shifted to stimulate real growth and contain undesirable disinflation (Fawley and Neely, 2013).

In general, when central banks opt for direct purchases of securities in the capital market, they have to set at least five key parameters to define their programmes. First, they need to decide on the asset class to be bought. Second, they are required to choose the respective volume in order to define the impact of the programme. Finally, with the selection of the remaining three parameters, the central bank can fine-tune the QE policy, particularly, in terms of maturity, rating and liquidity (Mertens, 2017). Over time, the ECB has implemented a wide range of purchase programmes varying these parameters to target specific risk premia. However, all of them share the common goal of easing funding conditions for financial and non-financial corporations. Table 1 summarizes the ECB's purchase programmes since 2009 in a chronological order.

The most significant shift in ECB's monetary policy has been the launch of the CSPP in March 2016, as a part of the APP. In contrast to previous purchases programmes, primarily focusing on the economy or the public sector in general, the key feature of the CSPP is that it is specifically directed at assets in the non-financial sector (Mertens, 2017). Basically, the CSPP represents a suitable alternative in providing credit to corporates by bypassing the banking system (Arce et al., 2017, Grosse-Rueschkamp et al., 2017). Despite the extra liquidity facilities provided by the ECB, banks have frequently been unable to adequately provide credit to the real economy due to, inter alia, the non-performing loan burden, higher regulatory requirements and ongoing restructuring. Faced with declining incomes, tightening prudential regulation⁷ and high levels of debt, banks rather have utilized the additional funds to deleverage their own debt positions (Joyce et al., 2012, Demertzis and Wolff, 2016).

As a consequence, the CSPP has been carried out - jointly with a further cut in the deposit facility rate and a new series of four targeted longer-term refinancing operations (TLTRO II) - to address these transmission failures. In this way the ECB can take on the risk that banks are currently unable or unwilling to take. Hence, the CSPP is to be understood as complementary to the main thrust of the APP supposed to enhance the impact of previous QE policies (European Central Bank, 2016, Abidi et al., 2017). Figure 2 displays the dynamics of the APP policy. While the Eurosystem's total holdings have increased constantly over time, the CSPP share over total APP holdings is comparatively small. In fact, the outstanding volume under the CSPP (€ 128 billion) accounts

solely for 6 percent of total APP holdings (€ 2,243 billion) by November 2017.

Addressing CSPP's effectiveness will almost inevitably be part of a bigger picture that includes insights into recent global monetary policy developments. Given this goal, it is useful at the outset to dig deeper into the technical features of the CSPP policy.

2.2. ECB's Corporate Sector Purchase Programme

The Governing Council of the ECB announced at its March 10, 2016 meeting the launch of the CSPP as an extended leg of its QE programme.⁸ Operations commenced three months later, on June 8th as demonstrated in Table 1. Under the CSPP, the Eurosystem buys debt securities issued by non-financial corporations with the goal of consolidating the pass-through of the monetary policy stance to the real economy. The Governing Council is committed to continue CSPP purchases without imposing any temporal restrictions. In conjunction with other non-standard measures in place, the CSPP is intended to stimulate spending and thereby maintain inflation rate levels below, but close to, two percent in the medium term (European Central Bank, 2016).

The programme is coordinated by the ECB, but carried out by six national central banks acting on behalf of the Eurosystem. These include the central banks of Belgium, Finland, France, Germany, Italy and Spain. Each central bank is responsible for purchases from issuers in a particular region of the euro area (Abidi et al., 2017). Hereby the ECB acts as a buy-and-hold investor. Assets purchased under the programme are held until maturity and the principal is reinvested even after a possible termination of the purchase programme (Grosse-Rueschkamp et al., 2017). As a necessary condition, assets must be acceptable as collateral for Eurosystem credit operations,⁹ subject to further criteria as explained hereafter (European Central Bank, 2016, Abidi et al., 2017). As mentioned previously, asset purchase programmes are carefully designed by central bank authorities to ensure that the bond portfolio purchased under the CSPP has a reasonable level of risk and a certain degree of diversification.

As part of its plan, the ECB buys only euro area bonds issued by non-financial corporations denominated in euro. The ultimate parent company may as well be located outside the eurozone region, however the issuer must be established within the euro area. Public undertakings and credit institutions that are subject to banking supervision are excluded altogether. This rule also applies to issuers that have any

⁸Note that the main technical parameters of the programme were announced on April 21, 2016 (Abidi et al., 2017).

⁹In compliance with its statute, the ECB provides credit only against adequate collateral. Collateral comprises both marketable and non-marketable assets. In general, assets that are accepted as collateral by the Eurosystem are labelled as 'eligible' and the eligibility is assessed by national central banks according to the criteria specified in the Eurosystem's General Documentation (Tamura and Tabakis, 2013). The list of marketable eligible collateral is updated daily and published on the ECB's website.

⁷Basel III has strengthened the capital adequacy and liquidity rules, upon which banks are required to adhere to.

Table 1: ECB's Purchase Programmes; Source: European Central Bank (2015), Andrade et al. (2016), European Central Bank (2017b), Mertens (2017).

The table lists all purchase programmes conducted by the ECB since 2009. The column 'End' states the termination date of the programme. As some purchase programmes are still in place, corresponding rows are yet empty. This is true for CBPP3, asset-backed securities purchase programme (ABSPP), PSPP and CSPP which are all part of the APP. The abbreviation p. m. in the Volume column indicates that purchases are carried out per month. Note that except for the CSPP all asset purchase programmes have been initiated to stabilize banks' balance sheets.

	Start	End	Asset Class	Volume	Rating	Maturity
CBPP1	July 2009	June 2010	Covered Bonds	60	AA	unlimited
SMP	May 2010	Sep 12	Government Bonds	unlimited	unlimited	unlimited
CBPP2	Nov 11	October 2012	Covered Bonds	40	BBB-	≤10.5y
CBPP3	October 2014		Covered Bonds	unlimited	BBB-	unlimited
ABSPP	Nov 14		Asset Backed Securities	unlimited	BBB-	unlimited
PSPP	March 2015		Public Sector Assets	30 p. m.	BBB-	2-30y
CSPP	June 2016		Corporate Bonds	80 p. m.	BBB-	6m-30y




Figure 2: ECB's Expanded Asset Purchase Programme Holdings; Source: ECB Statistical Data Warehouse (author's own computations).

The figure shows the Eurosystem's total holdings under the APP at the end of the month, represented by the right scale and denominated in euro. Moreover, the breakdown of the APP by each subprogramme is illustrated. The observation window ranges from October 2014 to November 2017.

parent undertaking which is a credit institution (Abidi et al., 2017, Grosse-Rueschkamp et al., 2017).¹⁰

Moreover, in order to qualify for purchase under the CSPP securities must satisfy a minimum credit rating of at least investment grade (BBB-/Baa3/BBBL) assigned by an external rating agency. In accordance with the practice followed under its collateral framework, the Eurosystem recognizes

¹⁰More precisely, bonds issued by an entity which is supervised under the Single Supervisory Mechanism are not eligible for purchase under the CSPP. At the same time, in order to ensure a level playing field between euro area and foreign issuers, issuers with a parent company that is subject to banking supervision outside the euro area are excluded as well (European Central Bank, 2016).

credit assessments by only four credit rating agencies such as Standard & Poor's (S&P), Moody's, Fitch Ratings and Dominion Bond Rating Services. Most noteworthy, in the event of a deterioration of the issuer's credit quality, the ECB is not obliged to sell its holdings (Grosse-Rueschkamp et al., 2017). Within this context, the first-best credit rating is relevant. More precisely, a bond that is rated below investment grade by three rating agencies except for one will still be eligible for admission to the CSPP programme (Abidi et al., 2017).

The maturity spectrum of debt securities can range between six months to less than 31 years at the date of the purchase. The upper bound is in line with that applied under the PSPP framework. The lower bound ensures that bonds

issued by small and medium-sized corporations are also included in the universe of qualifying debt instruments, while at the same time restraining the number of redemptions during the duration of the CSPP (European Central Bank Website, 2018a). In addition, the bond's yield to maturity has to exceed the level of the deposit facility rate at the time of the purchase (European Central Bank, 2016).¹¹

Further, as the ECB seeks a market-neutral implementation of the CSPP, purchases are conducted according to a benchmark defined at the issuer group level. The benchmark applied for purchases mirrors proportionally the market value of all eligible assets outstanding whereas the market capitalization serves as a weighting factor for the different jurisdictions within the benchmark. The purpose of issuer group level limits is to ensure a certain degree of diversification and neutrality in the allocation of purchases across corporations such that the overall portfolio is sufficiently heterogeneous (European Central Bank, 2017a, European Central Bank Website, 2018a). Total purchases under the CSPP should not exceed 70 percent of the issued value of each bond. At the same time there is no minimum issuance volume for eligible assets. This implies that bonds issued by small and medium-sized corporations with typically small issuance volumes can also be purchased (European Central Bank, 2016, European Central Bank Website, 2018a).

To sum up, there are seven conditions a bond must meet in order to qualify for CSPP purchases: eligibility as collateral for Eurosystem operations, a non-financial corporation issue, denomination in euro, an investment-grade rating, a yield of above the deposit facility rate, a maturity of between six months and 31 years, and an issue share limit of 70 percent per security.

In principal, under the CSPP the Eurosystem considers debt securities available in both the secondary and the primary market. In the latter case, it may participate in both public and private placements. In practical terms, these purchases take place concurrently and in competition with other investors, adhering to free-market principles (European Central Bank, 2016, Arce et al., 2017). The actual pace of purchases under the CSPP depends on prevalent market conditions. Monthly net purchases during the period from June 2016 to November 2017 have ranged between € 4 billion and € 10 billion. Overall, since the start of the programme in June 2016, on average corporate bonds worth € 7 billion have been bought monthly (see Figure 3). By the end of November 2017, CSPP purchases reached a total of € 128 billion, and were relatively diversified across ratings, sectors and countries. In general, CSPP holdings follow closely the CSPP-eligible bond universe, that is the composition of CSPP holdings mirrors that of the CSPP-eligible bond universe (European Central Bank, 2017a).

Although the ECB does not disclose the exact amount purchased for each bond, a recent ECB report states that, so far,

medium credit quality companies within the utility and consumer sector have attracted the most CSPP demand. The majority of purchases have been undertaken in the secondary market, with issues mainly from Germany and France (European Central Bank, 2017a).

2.3. Impact of Asset Purchases Programmes: Theory and Evidence

Over the last few years, given their widespread use by central banks, there has been a surge of theoretical and empirical research that aims to shed light on the workings of asset purchase programmes. The standard view in macroeconomic theory is that, in general, these programmes will not have any effects on the macro-economy, as the monetary policy stance is fully described by the current and expected future level of the nominal short-term interest rate. In line with this notion, QE policies are assumed to present a mere reallocation of assets from the balance sheet of private investors to the balance sheet of the central bank, while the reallocation as such does not change asset prices. The main assumption underlying this model is that of perfect substitutability of assets. A single representative and rational agent, subject to an infinite horizon and no credit restrictions, would then be indifferent between assets held by the central bank and her own assets (Eggertsson and Woodford, 2003, Woodford, 2012). Within this theoretical framework, the CSPP should therefore be ineffective.

In reality, however, neither financial markets are frictionless nor market participants do behave economically rational. Under these circumstances, the central bank may purchase significant quantities of assets in specific market segments hereof limiting the supply relative to the demand. Reducing the amount of bonds outstanding - by displacing some investors and reducing the holdings of others - will create a scarcity effect that arbitrageurs may not be able to eliminate. Given securities are not perfect substitutes, prices will rise and expected returns on the securities will fall, eventually suppressing the risk premia. Put differently, purchases will bid up the price of targeted assets thereby diminishing respective yields. As of the lower yields, the private sector will be incentivized to use the excess money in order to rebalance its portfolios. Private investors will demand assets that are similar in nature to the assets just sold to the central bank. Subsequently, the downward pressure on yields will not necessarily be limited to the particular asset type purchased but spill over to other asset classes in neighbouring markets (Vayanos and Vila, 2009, Bernanke, 2010). Eventually, the underlying mechanism of this portfolio balancing effect will also lead to lower interest rates relevant to consumption and investment spending. In fact, depressed yields imply lower borrowing costs for firms and households, which in turn will stimulate spending. In addition, higher asset prices enhance spending by the implicit increase in the net wealth of asset holders (Joyce et al., 2011a,b).

While the majority of empirical studies affirm the successful transmission through the portfolio balance effect, some researchers suggest the existence of novel channels that may

¹¹The deposit facility rate is the interest banks receive for depositing money with the central bank overnight (Kojien et al., 2016).




Figure 3: CSPP Monthly Net Purchases by Transaction Method; Source: ECB Data Statistical Warehouse data (author's own computations).

The figure shows the breakdown of primary and secondary market monthly net purchases under the CSPP, denominated in euro. The observation window ranges from June 2016 to November 2017.

be at work. For instance, recent literature has detected the signaling channel, through which asset purchases by monetary authorities may affect the economy. The mechanism of the signaling channel operates indirectly when market participants interpret and infer information from monetary policy announcements. Signals such as QE announcements may be viewed as a commitment by the central bank to keep expected short-term interest rates low for an extended period of time (Grosse-Rueschkamp et al., 2017). Alongside the signaling and portfolio rebalancing channel, Joyce et al. (2011b) address the liquidity and confidence channel. Similarly to Mertens (2017), they refer as well to the bank-funding channel aimed at increasing liquidity in the banking sector, which they, however, condemn as ineffective during times of severe financial crisis. Krishnamurthy and Vissing-Jorgensen (2011) have been able to make a pivotal contribution and extend the existent literature by launching in total seven channels through which unconventional monetary policy can transmit its effects. They discuss, for instance, the default risk channel that acts through reducing corporate default risk. If the CSPP indeed succeeds in stimulating the economy, it can be expected that the credit default risk of corporations will drop. Standard asset pricing models predict that investors' risk aversion will also fall as the economy recovers. More specifically, favourable market conditions are related to an increase in investors' risk appetite, underscoring the lower default risk perception, and ultimately implying a lower default risk premium (Fontana and Scheicher, 2016, Krishnamurthy and Vissing-Jorgensen, 2011). Gilchrist and Zakrajsek (2013) argue in the same manner but do not explicitly refer to the default risk channel. Furthermore, Eser and Schwaab (2016) agree that asset purchases can affect the default risk perceptions of market participants. However, they

claim that this is to be attributed to the signaling channel.¹²

Overall, lack of consistency among researchers with respect to the definition and the interpretation of channels makes it difficult to pinpoint the CSPP impact and link it to a single transmission channel. In my analysis I will follow recent remarks by the ECB policymaker Coeure (2017) who argues in favour of the standard portfolio rebalancing channel as the main transmission mechanism for the APP as a whole. Moreover, I will focus on the default risk channel which is by its nature particularly relevant for this study. The assessment of the remaining channels will not be the object of this study due to the limited scope of this paper.

While the exact transmission process is debated heavily in the literature, there is broad consensus about the effectiveness of unconventional monetary policy.¹³ Within this context, studying the impact of asset purchases on market prices provides the starting point for assessing a policy's effectiveness, as any QE intervention is very likely to have an impact directly on markets where purchases have been conducted, and indirectly on neighbouring markets. Hereafter, I aim to summarize the most relevant contributions with a special focus on the literature that deals particularly with the CSPP.

Looking at the US evidence, there is a large and growing body of literature that analyses QE policy effects on asset prices. Pioneering evidence is presented by Gagnon et al. (2011) in an event study on the Federal Reserve's purchases between December 2008 and March 2010, hereafter re-

¹²While Krishnamurthy and Vissing-Jorgensen (2011) differentiate between the signaling channel and the default risk channel, Eser and Schwaab (2016) relate default risk to the signaling channel.

¹³Note, however, that the estimated size of effects varies considerably across studies. Heterogeneity in results derives from different measuring methods.

ferred to as QE1. QE1 included a variety of assets such as mortgage-backed securities, treasury securities and agency securities; and proved to have economically meaningful and long-lasting effects on longer-term interest rates. Based on key QE1 announcements dates and time series regressions, Gagnon et al. (2011) notice large changes in the 10-year Treasury yield relative to the 2-year Treasury yield. To put this result into perspective, the QE1 policy has worked predominantly by mitigating the term premium. Indeed, the 10-year term premium was estimated to have been reduced between 30 and 100 basis points overall.

In consonance with the former study, Krishnamurthy and Vissing-Jorgensen (2011) target the effect of the Federal Reserve's QE1 programme through an event study methodology but dig deeper into the QE mechanisms. They find that this policy had a significant effect on yields, inter alia, through the default risk channel. In fact, they observe a substantial drop in nominal interest rates on lower-rated corporate bonds. Most strikingly, however, the authors report declining CDS rates linked to a clear pattern across credit ratings, ranging from Aaa to B. On the event dates related to QE1, there is a large decrease in CDS premia especially for lower grade firms. In particular, 5-year CDS rates of Aaa firms do not change appreciably with QE1 (6 basis points), whereas 5-year CDS rates written on B rated firms experience the largest fall (991 basis points). In terms of statistical significance, two-day changes in CDS spreads are significantly more negative for QE1 announcement days than on other days for 4 of 6 rating categories. Altogether their study suggests that - consistent with the default risk channel - QE has reduced the default risk premium.

Similarly, Gilchrist and Zakrajšek (2013) focus on the market's default risk perception and research the sensitivity of credit risk - measured by means of CDS indices - to changes in the benchmark market interest rates prompted by the US QE announcements. The authors apply a heteroscedasticity-based approach and find that the policy announcements have substantially lowered the overall level of credit risk in the economy. More specifically, for both the investment- and speculative-grade U.S. corporate sector there are economically large and statistically significant reductions in CDS index spreads. In line with the former study, the decline in the lower-rated CDS index is larger than in the higher-rated segment. In the financial sector, however, the response on credit risk is much more muted. A range of subsequent studies provide supportive findings in that the Federal Reserve's QE asset purchases were successfully diminishing medium and long-term interest rates, including those by Hancock and Passmore (2011), Swanson (2011), Hamilton and Wu (2012), Neely (2012), Wright (2012) and D'Amico and King (2013).

In order to avoid cultural bias and gain a sense of the universal challenges, it is crucial to investigate whether these trends appear in other countries as well. Undeniably, for the United Kingdom, Meier (2009) and Joyce et al. (2011a,b) find that the Bank of England's asset purchases between March 2009 and January 2010 had economically significant effects on government bond yields. Based on an event

study approach, Meier (2009) determines that the initial QE announcements have reduced government bond yields between 35 to 60 basis points. Joyce et al. (2011a) estimate that medium- to long-term government bond yields haven fallen cumulatively by around 100 basis points. In addition, they report a downward trend for corporate bond yields with smaller effects on investment grade bonds and larger effects on non-investment grade bonds. Further insights into the significant impact of the first phase of Bank of England's QE policies have been provided by Joyce and Tong (2012) and McLaren et al. (2014). For Japan there is also compelling evidence that outright asset purchases have led to a drop in long-term yields and a boost in asset prices (Lam, 2011, Ueda, 2012, Fukunaga et al., 2015).

For the euro area, there is a set of studies that qualitatively supports the results ascertained in the US, the United Kingdom and Japan. For instance, Andrade et al. (2016) scrutinize the impact of ECB's APP announcement across 24 studies to find a persistent decrease in 10-year sovereign yields with effects being the largest when new interventions are announced. Additionally, the researchers make efforts to take into account the banking sector. Particularly, the APP induces an increase in share prices of banks subject to a higher proportion of sovereign bonds in their portfolios. Apart from that the authors employ a general equilibrium model to compare the APP to conventional monetary policy measures. Hereof they argue that the APP has had an impact similar to a 100 basis point interest rate cut. Altavilla et al. (2015) obtain a similar set of findings confirming the economically meaningful impact of the APP on asset prices. They detect that the reduction in yields is more pronounced for longer dated sovereign bonds in high-yield countries. Interestingly Altavilla et al. (2015) document spillover effects. In particular, a decrease in euro area sovereign bond spreads by 100 basis points leads to a statistically significant decrease in corporate spreads by 63 basis points and 50 basis points for financial and non-financial institutions, respectively. The authors argue that this is to be attributed to the interplay between the transmission channel and the degree of financial distress. Similar patterns also show up in a succeeding study by De Santis (2016b) who accounts for the fact that the APP was implicitly communicated to the market before actual purchases had started. His econometric analysis suggests that the ECB policy has reduced GDP-weighted 10-year euro area sovereign yields by 63 basis points over the period from September 2014 to October 2015, with vulnerable countries benefiting the most.

Although the SMP is in general not considered a QE policy, in the literature there is some evidence that the programme works, inter alia, through the default risk channel and accounts for spillover effects. Hence, in the following I will briefly outline relevant contributions. Within the SMP framework, the ECB has engaged in purchases in five distinct sovereign markets beginning with Greece, Ireland, and Portugal and then expanding the programme to Spain and Italy. Based upon a panel regression model Eser and Schwaab (2016) evaluate the yield impact of the SMP in the euro area

sovereign bond market from 2010 to 2011. The authors estimate that government bond purchases have been successfully declining yields for Greece, Ireland, Italy, Portugal and Spain. For instance, in Greece € 1 billion of bond purchases have lowered yields by more than 20 basis points. Further, in their study Eser and Schwaab (2016) show that SMP purchases have also affected CDS spreads, yet to a lesser extent as compared to corresponding sovereign bond yields. While CDS spreads for Greece have reduced by 10 basis points, interestingly, for Italy the SMP impact on CDS has been positive. The authors conclude that a positive impact of purchases on CDS but not on the bond yield could be an indication of market participants worrying about moral hazard but welcoming the reduced liquidity risk premia on bonds. A related study by Kojien et al. (2016) estimates that the ECB is exposed to 3 percent of all sovereign risk as a consequence of the SMP intervention.

The latest literature urges to widen the research examining a set of ECB's asset purchase programmes conjointly. In a comprehensive study, Szczerbowicz et al. (2015) finds that SMP, OMT and CBPP have been effectively lowering refinancing costs of banks and governments, especially for periphery countries in the euro area. Further, she reports spillover effects to non-targeted asset classes, particularly, a 19 basis points tightening of covered bond spreads upon the SMP announcement and a 5 basis points tightening of sovereign bond spreads upon the CBPP announcement. In her study, she employs an event study approach based on daily data throughout the time period from 2007 to 2012. Transferred to a broader sample, Fratzscher et al. (2016) document reduced risk aversion, higher equity prices and lower credit risk for sovereigns and global banks upon the ECB intervention. Most noteworthy, as a consequence of the announcement of the OMT and SMP, equity prices have increased globally, while contraction in bond yields have been concentrated in periphery countries within the Eurozone. For Italy and Spain, for example, the 10-year government bond yield has declined cumulatively by 74 and 121 basis points, due to OMT and SMP related announcements, respectively. Consecutive work by Krishnamurthy et al. (2017) examines the relationship more closely proposing that both SMP and OMT have been much more effective at reducing sovereign bond yields than the LTRO measures across Italy, Spain and Portugal. Based on the Kalman-filter augmented event study, their analysis reveals that default risk accounts for 37 percentage of the reduction in yields. At this early stage, available empirical evidence on the CSPP is limited but points towards a similar direction as earlier QE studies. According to the ECB the announcement of the CSPP as such in March 2016 had a significant impact on the secondary market pricing of corporate bonds. Specifically, the 5-year yield (spread) on euro area CSPP-eligible bonds has decreased steadily in the period following the announcement. This downward movement is consistent across all credit rating classes although more pronounced for lower-rated bonds (European Central Bank, 2017a). Further empirical research supports the former findings. For example, Abidi et al. (2017) demonstrate

that the CSPP leads to a significant decrease in euro area corporate bond yield spreads by around 40 basis points. Contrary to expectations, they find that the decline is more noticeable in the sample of non-eligible bonds close to the investment grade threshold. In addition, the authors document an increase in bond issuance volume, in particular for non-eligible bonds. This is an important insight implying that - in line with the notion of spillover effects - the CSPP impact is not limited to eligible bonds but extends beyond the eligibility criteria.

Given favourable credit conditions induced by ECB's expansionary monetary policy, large corporations are increasingly able to finance themselves through bond issuances rather than bank loans. At the same time, the inception of the CSPP has deepened the corporate bond market with an expanded primary market activity. Arguing in this line the consequence would then be the contraction in the demand of bank loans as a funding source creating capacity in the balance sheet of banks. Consistent with the objectives set for the CSPP programme, banks should therefore be willing to divert the flow of credit towards companies that do not rely on capital markets for their financing, particularly small and medium-sized enterprises. Indeed, Arce et al. (2017) observe for Spanish companies that the CSPP has not only achieved its direct goal of reducing financing costs and stimulating new debt issuances but also has benefited non-eligible firms by means of a subsequent reallocation in the loan base of banks, especially in conjunction with the TLTRO. In relative terms, one euro less in the credit balance of eligible issuers leads to an increase of around 78 cents of euro in the credit balance of non-eligible firms following the CSPP. On a broader level, Grosse-Rueschkamp et al. (2017) take up this notion and validate that the intervention in the bond market indeed has reduced corporates' reliance on the banking system across the euro area, especially of investment grade corporates with lower credit quality. For their analysis, the authors use a more representative sample of publicly listed firms in S&P's Capital IQ. The DID framework in their study then reveals that within the set of CSPP-eligible firms, BBB rated firms increase their bond leverage relative to higher rated firms (1.6 percentage points versus 1.2 percentage points). Akin to the former study, Grosse-Rueschkamp et al. (2017) confirm that banks subject to a high proportion of CSPP-eligible firms in their portfolios prior to the CSPP announcement subsequently shift lending to private ineligible firms.

In conclusion, the impact of the CSPP programme spans two main dimensions: a relaxation of corporate lending costs and the spillover to non-targeted assets. Hereof empirical research considers predominantly the corporate bond market, while evidence on the derivative market is rather scarce.

3. Credit Default Swaps

CDS are classified as credit derivatives. These bilateral contracts present a relatively recent financial innovation. The first contract was traded by J.P Morgan in 1994 to meet the

increasing demand for transferring counterparty credit risk. Since then the market has grown remarkably (Augustin et al., 2014). Nowadays the CDS is the most popular and widely used instrument amid the broad class of credit derivatives, inter alia, due to its high degree of convenience with which market participants can express a view on the credit market (Blanco et al., 2005, Longstaff et al., 2005).

3.1. Overview of Corporate Credit Default Swaps

Single-name CDS are useful instruments to offset exposure to counterparty credit risk, namely the default risk of a certain issuer of debt capital. More precisely, two parties enter into an agreement, whereby the CDS buyer acquires protection from the CDS seller against the default of a third party, called the reference entity or the name. The reference entity, a particular company, can be either the issuer or the guarantor of the debt obligation. In essence, a CDS contract can be interpreted as an insurance, since one party intends to insure against the possibility of default while the other party is willing to bear this risk. Technically speaking, the protection seller 'longs' a third-party credit risk, whereas the protection buyer 'shorts' the credit risk (Blanco et al., 2005, Fontana and Scheicher, 2016). In contrast to a classical insurance contract, however, an engagement in a CDS does not require ownership of the reference asset. In effect, speculators are able to take long (short) positions in credit risk by selling (buying) protection without the need to trade the underlying bond (Blanco et al., 2005, Stulz, 2010, Breitenfellner and Wagner, 2012).¹⁴ Hence, investors who provide the capital are not necessarily those who bear the credit risk. This is an important insight as according to Stulz (2010) the separation of funding and risk bearing introduces greater transparency in the pricing of credit.

Back to a standard CDS contract, protection is sold in exchange for the payment of a regular fee at fixed payment dates (Fontana and Scheicher, 2016, Breitenfellner and Wagner, 2012). As in an interest rate swap (IRS) agreement, the fee is set such that the initial value of the CDS is zero which means there is no cash exchange at the time of trade.¹⁵ This fee is an annual premium paid over the lifetime of the contract, generally referred to as the CDS spread or CDS premium. It is denominated in percentage of the notional amount insured or in basis points, and to be paid in quarterly or semi-annual instalments (Augustin et al., 2014) until the maturity of the contract or the occurrence of issuer default (whichever comes first) (Fontana and Scheicher, 2016). Most importantly, a CDS contract is written on a single company rather than on specific bond issues (Chen et al., 2010). Hence a CDS usually comprises a category of the capital

structure, such as the senior, unsecured, or junior debt obligations of the underlying entity, and references a particular amount of the insured debt, defined as the notional amount (Augustin et al., 2014).

If a default does not occur over the lifetime of the contract, then the contract will expire at its maturity date and the protection seller will not pay any compensation. Conversely, in the case of a default, the contract is terminated prematurely and the protection component is triggered, which in fact is a cash payout reflecting the loss experienced by holders of defaulted debt obligations (Fontana and Scheicher, 2016, Breitenfellner and Wagner, 2012). The protection component is linked to a specific credit event. This contingent credit event refers to the case when the underlying entity fails to meet its obligations for any of a predetermined set of its debt claims, designated as the reference obligation. Formally, the occurrence of a credit event must be documented by public notice and notified to the investor by the protection buyer (Augustin et al., 2014). Amid the class of qualifying and valid credit events are bankruptcy, failure to pay, obligation default or acceleration, repudiation or moratorium (for sovereign entities) and restructuring, whereas the International Swaps and Derivatives Association (ISDA)¹⁶ eventually decides on whether a credit event has occurred. Put differently, credit events adhere to the strict standardised definitions laid down by the ISDA. For example, according to the ISDA documentation the restructuring event refers to the case when either the interest rate or the principal paid at maturity is reduced or postponed, a priority ranking of payments is altered, or when there is a change in the currency or composition of payments (O'Kane et al., 2003, Blanco et al., 2005, Beber et al., 2009).

Following a credit event, the final settlement can be cash or physical delivery, depending on the terms of the contract. Either the protection seller compensates the protection buyer for the incurred loss by paying the face value of the bond upon delivery of the defaulted bond (physical settlement), or by paying the difference between the postdefault market value of the bond and the notional value (cash settlement). In particular, with cash settlement the post-default value of the bond is determined through an auction mechanism. The monetary exchange involves then only the actual incurred losses while the protection buyer continues to hold on to the debt claim on the underlying reference entity's balance sheet, given she owns the claim (Fontana and Scheicher, 2016, Augustin et al., 2014).

While in the early days of CDS market participants had the choice of settling physically or in cash upon the occurrence of a valid credit event, for practical reasons, a marketwide cash settlement mechanism has been implemented in recent years. The main concern is that, with CDS outstanding greater by multiples than the volume of bonds issued, the bond market is subject to occasional market squeezes. Ef-

¹⁴Note that, as from November 2012, the European Union has enacted a regulation that bans short sales of uncovered sovereign debt CDS; corporate CDS are not subject to this regulation (Regulation, 2012).

¹⁵Both legs of a CDS need to have the same value at the inception of the swap. This is known as the zero net-present-value condition for swaps and implies that engagement in a CDS does not require a principal payment (Longstaff et al., 2005, Fontana and Scheicher, 2016).

¹⁶ISDA provides guidance on legal and institutional details of CDS contracts. The association has played a significant role in the growth of the CDS market by providing standardized contracts in 1992, the ISDA Master Agreement, which has been updated continually since then.

fectively, only few deliverable cash bonds are available in the market to settle all CDS trades (Blanco et al., 2005). Investors recognizing this are incentivized to source bonds, thereby raising artificially the bond price beyond the expected recovery value, and also increasing the volatility of the post-default bond. As a consequence, with the introduction of the Big Bang and Small Bang¹⁷ protocols, cash settlements have become gradually convention (Augustin et al., 2014).

As depicted in Figure 4, under standard physical settlement the protection buyer has to deliver a bond of seniority at least equal to the obligation referenced in the contractual agreement in the case of a default. In return, the buyer will receive the full notional amount of the underlying contract. The protection seller can then try to maximize the resale value of the debt claim received or continue to hold on to it. Most noteworthy, if the credit event occurs in between the regular premium payment dates, then at the final settlement the protection buyer will also have to pay that part of the premium to the protection seller that has accrued since the most recent payment (Longstaff et al., 2005).

Against this backdrop the restructuring event is an interesting feature of CDS contracts. It is considered a 'soft' event because, in stark contrast to other credit events, it allows for debt restructuring prior to any violation of the contract. More specifically, provided a firm is in financial distress but still economically viable, it may be optimal for the firm to restructure its debt within a private or debt workout while continuing operations. Within the context of a physical settlement, naturally, some deliverable reference bonds will be cheaper than others, such as debt with long maturities and low coupon rates. If there are multiple bonds available for delivery, the protection buyer will most likely choose to transfer the 'cheapest' bond to the protection seller.¹⁸ Hereof restructuring clauses constrain the set of bonds that are available for delivery upon the occurrence of a restructuring event, and specifically prevent the delivery of very long-dated bonds. In general, there are four different types of restructuring events: the old restructuring clause, the deletion of restructuring as a credit event, the (American) modified restructuring and the (European) modified-modified restructuring.¹⁹ Intuitively, these restrictions reduce the value of the cheapest-to-deliver option, and in turn are an important determinant for the pricing of CDS. The higher the value of the

inherent option to the protection buyer is, the higher the restructuring premium and correspondingly the CDS premium will be (Blanco et al., 2005, Longstaff et al., 2005, Berndt et al., 2007, Augustin et al., 2014). Contractual clauses attached to the different restructuring credit events have been adjusted and updated several times by the ISDA. The recent 2014 definitions introduce a number of simplifications to the Big Bang and Small Bang protocols (Augustin et al., 2014). Overall, while the restructuring event is particularly relevant for the cheapest-to-deliver option, it represents at the same time the most critical aspect in the pricing of CDS contracts.

CDS are the most popular and widely used instrument amid the broad class of credit derivatives. On the one hand, CDS allow the mitigation of counterparty risk exposure, especially for capital or credit exposure constrained businesses such as banks, pension funds or insurance companies (Longstaff et al., 2005, Abad et al., 2016). For instance, CDS are often used by banks for risk management purposes and are recognized by regulators as a regulatory capital relief (Augustin et al., 2014). On the other hand, speculation is a significant driver for engagement in the CDS market. Besides hedging investors are able to gain speculative benefits, specifically from negative credit events. For example, investors buy CDS not necessarily because they expect a default but because they anticipate that CDS spreads will increase further. To cash in the profits, investors will not be obliged to wait for a default but can rather sell another CDS.²⁰ To sum up, CDS allow pessimistic investors to bet against prices (Delatte et al., 2012). Similar to other derivatives, CDS can be viewed as 'side bets' on the underlying assets without any effect on the fundamentals of these assets (Liu et al., 2017).

Generally speaking, CDS are over-the-counter transactions, not traded on an organized exchange, whereby trades usually take place between institutional investors and dealers (Longstaff et al., 2005, Augustin et al., 2014). While dealers assume the intermediary role, financial institutions, including hedge funds and mutual funds, non-financial corporations, as well as insurances and pension funds are net buyers of protection. The market is highly concentrated among a few dominant dealers with the majority of trades relating to a few reference entities (Augustin et al., 2014, Abad et al., 2016). In principal, CDS can be negotiated at any time and in unlimited amounts (Delatte et al., 2012). However, as a necessary condition, institutional investors and dealers have to enter into an ISDA Master Agreement, setting up the legal framework for trading. The ISDA Master Agreement specifies the contractual terms and provides investors with a fully documented yet flexible contract (Augustin et al., 2014).

3.2. Credit Default Swaps as a Measure of Credit Risk

Credit or default risk associated with a particular company can be quantified by a number of metrics. Tradition-

¹⁷The landscape for CDS altered significantly with the implementation of the CDS Big Bang and CDS Small Bang protocols on April 8, and June 20, 2009 for the American and European CDS markets, respectively. The primary goal of these market changes - mainly affecting the contract and trading conventions - was to improve the efficiency and transparency within the CDS market.

¹⁸Conceptually, this cheapest-to-deliver option is equivalent to a short position in a put option. If not otherwise specified in the contract, upon exercise the protection buyer will have the right to deliver the least valuable asset among the defined set of eligible reference obligations as long as they rank pari passu with the reference asset (Blanco et al., 2005). For empirical evidence on the cheapest-to-deliver option inherent in corporate CDS see Jankowitsch et al. (2008).

¹⁹For an in-depth discussion of the restructuring feature, see O'Kane et al. (2003) or Berndt et al. (2007).

²⁰Given an investor wants to liquidate her CDS position, it is more convenient to simply enter into a new swap in the opposite direction than trying to sell the current position (Longstaff et al., 2005).




Figure 4: CDS Transactions under Physical Settlement; Source: [Markit Group \(2008\)](#).

The figure describes CDS transactions under physical settlement. Hereby the protection buyer makes fixed periodic payments to the protection seller, for instance on a quarterly basis. Given a default event occurs, a payout is triggered. The protection buyer transfers the obligation referenced in the contractual agreement - not necessarily the defaulted bond - and in return receives the full notional amount of the underlying. In other words, the protection seller is obliged to buy back the defaulted bond at par value.

ally, the financial health of companies has been assessed by predicting default probabilities. These probabilities have often been derived by modelling historical default events in a logistic-regression framework or by applications of Merton's firm value model.²¹ However, in real life very few companies do default. As such, these frameworks are hard to calibrate empirically and subject to the rare-event bias ([Opsahl and Newton, 2015](#)). To overcome potential biases, I will follow [Krishnamurthy and Vissing-Jorgensen \(2011\)](#) and choose an outcome variable that is quantified for a relatively large number of companies as well as closely related to the risk profile of a company, namely CDS contracts written on a particular company.²² Given the underlying company becomes more risky, respective CDS rates will increase and vice versa. Accordingly, CDS are considered reliable measures of a firm's credit quality, widely used by practitioners and academics to gauge the market's perceptions of a firm's credit risk. In the broader sense, CDS spreads may also serve as a proxy for the firm's cost of wholesale funding ([Beau et al., 2014](#)).²³

²¹Such structural default models model explicitly the link between equity and default ([Merton, 1974](#)).

²²In particular, [Krishnamurthy and Vissing-Jorgensen \(2011\)](#) use CDS to isolate default risk premium effects for their estimation purposes.

²³There exists a link between the quality of borrowers' balance sheets and their access to external finance. Given profits decline and balance sheets

It may well be argued that within financial markets there are several alternative parameters that can be used to measure credit risk. Scanning the market for instruments with near-identical risk and return characteristics as a CDS, while abstracting from arbitrage, enables the identification of these parameters. Conceptually, in an arbitrage-free market a CDS could be replicated by an asset swap, which is a combination of an IRS and a defaultable coupon bond. The IRS swaps the coupon of the bond into a reference rate plus spread. The asset swap is chosen such that the value of the whole package is par value of the defaultable bond. However, the arbitrage is not perfect. Unlike CDS, IRS are not affected by credit events and thus not automatically cancelled at default ([Duffie, 1999](#)). Therefore, spreads of both, the asset swap and the CDS, can trade at different levels in the market for the same issuer and maturity. This differential in spreads is called basis. [Skinner and Townend \(2002\)](#) claim that CDS contracts resemble American put options on the underlying bonds. This is most evident under physical set-

deteriorate, bond investors will anticipate that the expected future cashflows will not meet the current debt obligations. In turn, as they will have to assume the additional credit risk, investors will demand a higher credit risk premium which subsequently increases the external cost of funding, and vice versa ([De Santis, 2016a](#)).

tlement when the underlying asset is delivered upon exercise. Duffie (1999) provides a more precise theoretical relation and argues that in absence of arbitrage opportunities the CDS is identical to a swap of a default-free floating rate note for a defaultable floating rate note. Hull and White (2000) build on this pricing model as a key element for the valuation of CDS.

Despite each of these instruments representing a theoretically legitimate measure of credit risk, in reality corresponding spreads across these instruments are not at parity in the short-run for reasons related to liquidity, margin requirements or simply market frictions (Fontana and Scheicher, 2016, Augustin et al., 2014). In this regard, Blanco et al. (2005) conclude that the CDS rate provides rather an upper limit on the price of credit risk. Overall, the choice of CDS as the variable of interest - in comparison to similar financial instruments - still is preferable due to simplicity and data availability.²⁴

Nonetheless, except for a few papers, my focus on CDS presents a clear deviation from standard literature that assesses the CSPP impact preliminarily with respect to bonds. At first glance, as the bond market is the predominant area targeted, the current approach seems reasonable, especially given the fact that in theory CDS and bond spreads are closely interlinked and should therefore provide equivalent results. However, it has been detected empirically that CDS spreads portray corporate credit risk better than corporate bond spreads. In other words, the CDS market clearly dominates the bond market in terms of modelling credit risk. This line of argumentation rests on two main pillars. First, in CDS markets pure issuer credit risk is priced.²⁵ After all, in absence of market frictions, the price of a CDS is solely about the expected default loss and not affected by contractual provisions such callability, maturity or coupon. In the bond market, in contrast, issue-specific credit risk and market risk are priced in a bundle (Norden and Weber, 2009, Fontana and Scheicher, 2016, Stulz, 2010). Second, price discovery takes place predominantly in the CDS market, that is default-risk related information is reflected earlier in the CDS market. While the first argument is debated heavily in the literature, the observed empirical difference between CDS and bond spreads is indeed proven to be due to informational problems and market frictions. To shed light on this matter, in the following I will briefly review the literature on the different dynamics of cash and derivative markets.

Pioneering work by Longstaff et al. (2005) is fielded using CDS data of 68 US firms from March 2001 to October 2002 to examine weekly lead-lag relationships between CDS spread changes, corporate bond spreads and stock returns.

²⁴In practice, for instance, spreads on corporate par yield floaters are difficult to observe (Hull and White, 2000).

²⁵Yet other studies dispute the validity of the underlying notion, arguing that CDS rates are not a pure measure of default risk after all, since they also incorporate a liquidity component (Fulop and Lescourret, 2007, Tang and Yan, 2007). Moreover, Jarrow (2012) discusses problems with using CDS to infer implied default probabilities on firms or sovereigns.

In their analysis the authors utilize CDS as a tool to disentangle default from liquidity risk in corporate bond spreads, as they assume that illiquidity is the non-default component affecting bonds but not CDS. Indeed, they find that information flows first into stock and credit derivative markets and then into corporate bond markets. Yet, their study shows no clear lead of the stock market over the CDS market, and vice versa.²⁶

Blanco et al. (2005) explore the same relationship but suggest in contrast to the former study that credit risk in CDS and bond markets is priced relatively equally. In cases where there is a deviation between corporate bond spreads and CDS premia, they attribute the difference to the tendency of CDS premia to lead corporate bond spreads in price discovery. Besides the authors argue that only well informed investors trade in CDS markets. Their dataset includes a daily time series for 33 U.S. and European investment grade companies during the period from January 2001 to June 2002.

In similar fashion, but based on a longer sample period, Zhu (2006) attests for a set of 24 investment grade firms from 1999 to 2002 that the CDS market leads the bond market in terms of price discovery. According to this study, CDS and corporate bonds spreads from the same firm with the same maturity horizon are cointegrated that is they may considerably deviate from each other in the short-run but are strongly linked in the long run. The author concludes that this deviation stems from the higher responsiveness of CDS premia to changes in credit conditions. In a sample of 58 firms across US, Europe and Asia covering the period 2000 to 2002, Norden and Weber (2009) examine monthly, weekly and daily lead-lag relationships in a vector autoregressive model and further highlight the existence of a cointegration relationship.

Finally, Delatte et al. (2012) abstract from the linear price discovery model often used in the standard literature and propose a non-linear method. However, their database relies not on corporate but on sovereign CDS premia from developed member states of the European Union. Their results suggest that price discovery varies with the degree of market distress. In particular, only during periods of relatively high distress does the CDS market dominate the information transmission between CDS and bond markets. Liu et al. (2017) further confirm that the information revelation role of CDS is especially apparent when there is a negative information shock. Additional empirical evidence on the concept of price discovery with respect to CDS is documented by Acharya and Johnson (2007), Berndt and Ostrovnaya (2014) and Batta et al. (2016).

Altogether, the CDS market is informationally more efficient absorbing information at a faster pace. This superiority of CDS over bonds encompasses many aspects but most importantly has its roots in the synthetic nature of CDS which facilitates a continuous flow of transactions. For example,

²⁶More recent studies actually provide evidence that the equity market leads both the CDS and bond market, see for instance Forte and Pena (2009) or Hilscher et al. (2015).

CDS offer a convenient way to short bonds, whereas establishing a short position in the bond market is rather problematic (Norden and Weber, 2009). Especially in times of financial turmoil, when short sales are particularly valuable agents tend to retract from the underlying bond market (Delatte et al., 2012). In this regard, many economists argue that the existence of short sales, as such, makes a market more responsive to new information (Stulz, 2010).²⁷ At the same time, CDS are more flexible and less capital-intensive because they require no principal payments. In contrast, within the underlying bond market the purchase of a bond generates a large cash outflow at the initiation of the trade (Norden and Weber, 2009).

Moreover, bond spreads in the secondary market depend on the availability and specificity of the total amount of bonds outstanding, which in turn is related to the issuance activity of the single firm. Given investors buy bonds with the motive to hold them until maturity, this curbs market liquidity. Poor liquidity in the secondary bond market will then make the purchase of large amounts of credit risk difficult and costly (Blanco et al., 2005, Longstaff et al., 2005). In stark contrast, the CDS market is more standardised and less dependent on primary market issuances (Blanco et al., 2005, Norden and Weber, 2009). In fact, CDS can be negotiated at any time and in arbitrarily large amounts. And indeed, the CDS market has experienced extraordinary growth over the past years with CDS outstanding greater by multiples than the volume of bonds issued.²⁸ To conclude, sensitivity to liquidity effects reduces the ability of the bond market to reflect information as timely as the CDS market, especially in the short run. Further differences between CDS and bonds can emerge due to accrued interest, the cheapest-to-deliver option and/or counterparty risk (Delatte et al., 2012).

Although the derivative and the cash market can differ on the same maturity-same reference entity in the short-run, CDS and bonds still provide roughly contemporaneous information. This is most evident when taking a step back and recognizing that an investor can conduct a risk-free strategy by combining the purchase of a bond with the corresponding CDS (Chen et al., 2010, Fontana and Scheicher, 2016). This insight is particularly relevant for the remainder of the present paper. More specifically, the negative effect of the CSPP on bond spreads - thoroughly discussed in the literature - can be transferred to the CDS market and serve as an anchor to determine the direction of the CSPP effect on CDS spreads. In fact, the next section takes up this line of reasoning to form two hypotheses.

Taken together, the standardized documentation, the liquidity, the ability to customize terms, and the 'pure' credit focus makes CDS contracts convenient to express a view on the credit market, particularly in regard to the deterioration or improvement of a firm's credit quality. Hereby the CDS

spread represents the price, market participants are willing to pay, in order to offset exposure to the reference entity's default risk. Therefore, these market-based indicators can be viewed as an appropriate metric to isolate and quantify credit risk.

4. Data and Empirical Strategy

Now that I have highlighted the institutional features of the CSPP initiative and granted a brief overview of corporate CDS, I proceed by elaborating on the underlying assumptions, the estimation method and the sample data used for the estimation. The aim of this section is to provide a strategy that isolates the direct and indirect effects of the CSPP programme on those firms whose bonds have been eligible by the programme.

4.1. Hypotheses

According to recent literature, the announcement of the CSPP was successfully followed by a significant decline in the spreads of bonds issued by non-financial corporations (European Central Bank, 2016). This contraction in credit default risk - proxied by the bond yield spread or asset swap spread - establishes the main rationale for my first hypothesis, whereas I resort to CDS spreads as an alternative metric to quantify credit risk. If the transmission mechanism of the CSPP is the default risk channel, I will expect the following:

H0: CDS spreads for CSPP companies and non-CSPP companies decline around the CSPP purchase shock.

My first hypothesis claims that CSPP purchases have evoked a specific market reaction, particularly a contraction in CDS rates. Given the CSPP succeeds in stimulating the economy, one should observe a reduction in expected defaults and, as a result, a decline in corporate default risk. The implicit assumption is that the policy lowers bond yields in order to increase the expected repayments of bondholders. Standard asset pricing models predict that investors' risk aversion will also fall as the economy recovers. More specifically, diminishing CDS premia would then be related to an increase in investors' risk appetite, underlining the lower default risk perception, and ultimately the lower default risk premium (Fontana and Scheicher, 2016, Krishnamurthy and Vissing-Jorgensen, 2011, Gilchrist and Zakrajšek, 2013).²⁹ Thus, evidence in favour of H0 would be consistent with the programme's objective of lowering risk premia across the European non-financial corporate sector.

²⁹Default risk premium may also diminish due to the possibility of risk mitigation by means of CDS. Put differently, the CSPP effect will be corroborated, given the reduction of CDS spreads allows a firm's creditors to hedge their credit risk at a relatively lower cost. In turn, creditors' willingness to supply credit to the same CDS-referenced firm will increase. This is, however, not the object of this paper and leaves space for future research.

²⁷For a formal model, see Diamond and Verrecchia (1987).

²⁸As already mentioned, CDS are also affected by illiquidity, yet to a lesser degree than bonds. In particular, lack of liquidity is more pronounced for larger companies as compared to smaller companies (Stulz, 2010).

Figure 5 plots the evolution of CDS prices from 2015 to 2017 for both entities that have issued CSPP bonds and entities that have not issued CSPP bonds, hence CSPP companies and non-CSPP companies.³⁰ Clearly, asset purchase programmes have been launched in response to widening credit spreads reflecting the overall adverse economic developments in the Eurozone. At the aggregate level, this graphical evidence underlines my hypothesis that the CSPP has moved the credit market. In fact, following the announcement in March 2016, denoted by the left vertical line in the graph, decreasing CDS spreads are visible over the long run. Nonetheless, this can only be taken as tentative evidence supporting H0 as the announcement date and the subsequent decrease in CDS spreads may be driven by some latent omitted variables. In order to formally test whether the reduction in CDS rates is certainly caused by the new policy, I propose the DID estimation as elaborated in the next subsection.

On the contrary, the impact of the official implementation date of the CSPP, denoted by the right vertical line in the graph, is rather ambiguous with a slight pick-up shortly after June 8, 2016. In line with the efficient market hypothesis it may be argued that the announcement per se absorbs available price information immediately for all bonds at the aggregate level such that the implementation date on its own becomes trivial (Fama, 1970). In this respect, Arce et al. (2017) disclose in their study that the CSPP effect on bond yields is more attenuated for the implementation date as compared to the announcement date (7.6 basis points versus 46 basis points). Further, they report that during the first month of purchases the effects are slightly higher with a value of around 8 basis points. Taking this into consideration, it is reasonable to focus on the announcement effect of the newly implemented policy as a basis to derive the hypotheses. This is also the current practice in the literature (see for example Gagnon et al., 2011, Krishnamurthy and Vissing-Jorgensen, 2011, Arce et al., 2017). However, if anything, I expect a lower bound estimate on the CSPP effect.

In the context of the second conjecture, my paper is closely related to the work by Abidi et al. (2017). They document that the CSPP impact on bond yield reduction is most noticeable in the sample of bonds that have not been subject to CSPP purchases. Though at first glance this may appear counter-intuitive, a closer look suggests that higher credit risk firms - typically a subset of non-CSPP companies - are supposed to benefit from the new policy on a larger scale. Indeed, Krishnamurthy and Vissing-Jorgensen (2011) and Gilchrist and Zakrajšek (2013) detect the pattern that the decline in CDS rates, following a QE policy, is more profound for firms with lower credit quality. This line of argument

rests on the fact that benefits associated with the CSPP do not accrue selectively but extend to non-targeted assets.

Figure 5 allows a comparison of CDS premia between CSPP and non-CSPP companies. Not surprisingly, there is a high degree of comovement in the CDS spreads of these two groups, reflecting the exposure to common macroeconomic factors. Over the whole sample period, though, spreads for the CSPP group are on average lower than that of the benchmark. Given the strict eligibility criterion for CSPP purchases, such as preliminarily targeting investment grade bonds, it is not surprising that the ECB is more inclined to buy bonds associated with lower credit risk. Throughout the year 2015, the spread for both groups widens substantially, reaching its peak in early March 2016. Shortly before the announcement of the programme on March 10, 2016, spreads exhibit a considerable decrease. This fall in spreads continues around the date of the announcement and thereafter, interrupted only by temporary phases of uncertainty in May and June. The United Kingdom's referendum on the European Union membership may be related to widening spreads, but the effect seems short-lived (European Central Bank, 2016). Over the course of the second half of the year 2016, spreads decline more gradually. Overall, by November 2017, the CSPP group reaches a new all-time low of about 60 basis points, which marks a reduction of 80 basis points relative to the peak in early March 2016. The downward trend is, however, more pronounced for the non-CSPP group with a tightening in spreads by about 170 basis points, from roughly 270 basis points in March 2016 to around 100 basis points at the end of year 2017. Interestingly, from mid-2017 onwards, spreads of the CSPP and the non-CSPP group slightly converge. As of this date, spreads are also more stable. At the aggregate level, this may suggest that while credit risk has reduced overall, the impact on the non-CSPP group will be more striking. The existence of potential spillover effects dictates my second hypothesis which reads:

H1: CDS spreads for non-CSPP companies will decline relatively more than for CSPP companies around the CSPP purchase shock.

H1 supports the view that the transmission mechanism of the CSPP operates as desired beyond the eligibility criteria. In fact, the reduction in funding costs induced by the CSPP should be reflected in substantially lower costs of default insurance, especially for riskier credits. When CDS rates decline relatively more for non-CSPP firms as expected, the CSPP policy contributes endogenously through spillover effects, in line with the theory of the portfolio rebalancing channel (Altavilla et al., 2015).

4.2. Empirical Strategy

This subsection elaborates on the underlying estimation method to test the hypotheses formed in the previous subsection. A simplistic approach to estimate the impact of CSPP may be to compare CDS rates across entities issuing purchased bonds and non-purchased bonds while exclusively

³⁰The data is extracted from Markit and comprises all non-financial CDS outstanding, irrespective of the fact whether firms have issued bonds purchased by the ECB under the CSPP policy. Bonds and companies that are subject to the CSPP policy are labelled as CSPP bonds and CSPP companies, respectively. This applies analogously for non-CSPP bonds and non-CSPP companies. Note, however, that the definition of 'non-CSPP' here deviates from the definition provided in the Section 4.2 and the definition used in the empirical estimation.



Figure 5: CDS Spread Evolution by CSPP Purchases; Source: Markit (author's own computations).

The figure presents the evolution of weekly CDS prices. The CDS set is split among those firms issuing bonds that are effectively purchased under the CSPP (CSPP) and those firms issuing bonds not subject to the new policy (non-CSPP). The observation window ranges from January 2015 to November 2017. The CSPP announcement in March 2016 is denoted by the left vertical line, while the right vertical line labels the official start of the CSPP in June 2016.

focusing on the post-intervention time, that is the period succeeding the purchase date. However, biased estimates may result when prior to that period differences in prices between the two groups exist. In order to establish causality I adopt a DID analysis which explicitly considers differences in prices prior to the policy implementation. In particular, the DID method identifies causal effects by contrasting the change in outcomes pre- and post-intervention, for the treatment and control group.³¹ DID assumes that, in the absence of treatment, prices remain unchanged and then trends within treatment and control groups are equivalent. This assumption of parallel trends allows the averages of the time-invariant unobserved variables to differ between treated and control groups, provided their effects do not change over time (Bertrand et al., 2004, Lechner, 2011, Morris et al., 2013).³²

The basic DID regression is given by the following equation:

$$Y_{it} = \alpha + \beta \text{Treat}_{it} \text{Pot}_t + \gamma \text{Treat}_{it} + \delta \text{Post}_t + \epsilon_{it} \quad (1)$$

where i and t index firm and time observations. Y_{it} is

the outcome of interest, Treat_{it} (=1 if bond of firm i is purchased at day t) is a dummy for CSPP entities and Post_t (=1 beginning from the initial purchase date and thereafter) is an indicator for the post-CSPP period. Furthermore, the equation includes a constant α and a random error term ϵ_{it} . The coefficient β is the DID estimator and identifies the treatment effect of the CSPP, as the treatment is $\text{Treat}_{it} \text{Post}_t$. The evaluation period is seven trading days before and after the purchase date of the respective bond.

Controlling for fixed effects rules out the concern that findings are explained by heterogeneous effects of the CSPP policy on CDS rates. By this means I am able to capture firm specific differences such as unobserved differences in local economic environments, management quality, or the cost of capital (Gormley and Matsa, 2013). Thus, in my estimations the basic DID regression is complemented by the inclusion of fixed effects to control for unobservable time-invariant factors at the firm level as well as time-varying fixed effects. The baseline regression reads:

$$\text{CDS Spread}_{it} = \alpha + \beta \text{Treat}_{it} \text{Pot}_t + \theta_i + \vartheta_t + \epsilon_{it} \quad (2)$$

where θ_i are firm fixed effects and ϑ_t are day fixed effects. Moreover, CDS Spread_{it} is the outcome of interest. Note that Pot_t will be absorbed by the time fixed effects. But what if the relationship between CDS premia and the regressor is non-linear? As Figure 6 shows, the distribution of CDS spreads

³¹For an early study in this vein, see Ashenfelter (1978).

³²It is worth emphasizing that, in contrast to the current practice in the literature, I refrain from employing an event study methodology. The underlying rationale is that there are serious identification issues with this econometric approach, such as neglecting key announcement dates or ignoring the simultaneous implementation of policies (see for example Gilchrist and Zakrajšek, 2013 or Fratzscher et al., 2016).

is skewed to the left³³ whereas its logged value appears normally distributed.

When assuming a linear model, I may obtain biased estimates of the effects of the CSPP on spreads. That is why I adjust the benchmark regressions to a semi log estimation by changing the dependent variable to $\log \text{CDS Spread}_{it}$.

$$\log \text{CDS Spread}_{it} = \alpha + \beta \text{Treat}_{it} \text{Post}_t + \theta_i + \vartheta_t + \epsilon_{it} \quad (3)$$

All coefficients in the non-linear regression are unchanged except for the outcome of interest, that is $\log \text{CDS Spread}_{it}$. In all specifications, I cluster standard errors at the firm level. Apart from that, the regressions specified previously are estimated on daily CDS spreads using the DID approach in combination with robust standard errors to account for heteroscedasticity. Overall, my research hypotheses suggest decreasing CDS spreads for entities in both the control and treatment sample, and increasing spreads for entities in the treatment group as compared to the control group. Thus, support for the hypotheses requires $\beta > 0$ and $\delta < 0$.³⁴

Studying the CSPP impact empirically requires solving an identification problem which relates to the endogeneity of CSPP-eligible bonds. Indeed, given the nature by how the CSPP policy is implemented, bonds that are accepted in the CSPP portfolio differ systematically from bonds which are not. In fact, the assignment of bonds to the CSPP portfolio is discrete and follows strictly the specific eligibility criteria developed by the ECB (see Section 2.2). Therefore, the comparison between eligible and non-eligible bonds or firms is likely to capture the effect of these (observable and unobservable) differences rather than capturing the causal effect of the CSPP. For example, low credit risk issuers - most likely to be part of the CSPP due to fulfilling the admission criteria - might have been less credit constrained before the start of the CSPP, relative to high risk underperforming issuers (Grosse-Rueschkamp et al., 2017). Ignoring this issue would then lead to the underestimation of CDS rates, and in turn, to biased estimates based on the standard regression analysis.

I overcome this obstacle by restricting the sample to CSPP firms, thus a subset of firms that issues at least one bond that is eventually to be purchased by the ECB. At the same time by exploiting the gradual implementation of the CSPP policy, namely the time dimension with which bond purchases have been executed, a potentially exogenous source of variation is generated that I can use to estimate the effect of the CSPP reform. In other words, the distinct purchase dates allow me to compare a subsample of firms transferred primarily to

the CSPP portfolio (treatment group) with firms transferred later (control group).³⁵ This suggests that in the control sample there will be companies having issued at least one bond purchased under CSPP over the course of time since the basis for the selection of these companies is the bond purchase list published by the ECB (see next subsection). However, as the relevant purchase date is succeeding the post-CSPP period of seven trading days³⁶ for my estimation purposes these firms are not viewed as treated firms. In this light in the following, firms in the treatment and control sample will be referred to as CSPP and non-CSPP firms, respectively. Figure 7 schematically summarizes the empirical strategy.

The following example illustrates the empirical strategy: data is retrieved on a daily basis with the sample period beginning in $t = 0$. Firm A has issued a bond that is effectively purchased under the CSPP on day $t = 1$. Firm A will be assigned to the treatment group for the purchase day $t = 1$. Firm B from the same industry has issued a bond that is purchased on day $t = 8$, that is one week later. As this date is beyond the evaluation period, firm B serves as a control variable. In aggregate, the within CSPP-sample analysis mitigates concerns related to omitted variables.

Another challenging task is related to the fact that the ECB decides to purchase bonds on distinct days. While on the one hand this is convenient for the purpose of demonstrating a direct effect between the CSPP and the change in CDS spreads, on the other hand running a single regression on the full sample to pinpoint the aggregate effect of the policy on a single day will not be possible. Moreover, as elaborated earlier, the main assumption upon which the DID approach rests is that of parallel trends. However, clustering all firms together can violate the former assumption as firms across different industries may not be comparable. For example, the insurance sector is likely to be negatively affected by accommodative monetary policies, because the ability to generate adequate interest income is severely impaired when credit risk is low at the aggregate level (Mertens, 2017). From a statistical perspective, an industry specific analysis addresses the heterogeneity concern that may remain in each sub-panel. Hence I run the DID regression at the industry level for each of the industries identified. Observations within an industry context also allow adherence to the underlying assumption of the DID approach.

4.3. Sample Data

The construction of the sample data is constrained by the availability of CDS. Indeed, reliable CDS data is available for a very low number of companies. Against this backdrop, the data collection is separated into two parts. First I seek to collect CSPP sample data at the aggregate level from the ECB

³³Provided that the trade takes place between institutional investors and dealers, a left-skewed distribution indicates that most dealers exhibit low counterparty risk while a few dealers have higher counterparty risk (Giglio, 2014).

³⁴This in line with the work by Abidi et al. (2017). Note that a negative estimate for β would be at odds with the second hypothesis but by no means rule out the existence of spillover effects. Rather this scenario would indicate that the CSPP impact on non-targeted firms is not as strong as expected.

³⁵Abidi et al. (2017) rule out the issue of selection bias by only considering bonds close to the eligibility threshold. They assume that in this case the admission to the CSPP portfolio will be random. Grosse-Rueschkamp et al. (2017), in contrast, define non-rated European firms with public debt as the benchmark.

³⁶Purchases are published on a weekly basis, hence there is at least one week between each purchase date.




Figure 6: Distribution of CDS Spreads; Source: Markit (author’s own computations).

The figure shows the distribution of CDS spreads for the entire sample. The mean spread level throughout the entire sample equals 79.368 basis points and the median value is 64.706 basis points. The standard deviation of spread levels is fairly high at 58.57 basis points. Overall, the distribution is highly left-skewed but conforms closer to the pattern of the Gaussian normal distribution after transforming it with the natural logarithm.




Figure 7: Empirical Strategy; Source: author’s own contribution.

The figure shows that in both the treatment and the control sample there will be CSPP companies in the course of time. The assignment to treatment remains reasonable as it is based on bonds purchased the first time through the CSPP. At the same time, companies in the control sample will have issued CSPP bonds eventually. However, the respective purchase date is succeeding the post-evaluation period of seven trading days.

website; subsequently I will combine the dataset with CDS pricing data extracted from Markit. The peculiar admission to the treatment and control sample as proposed in the pre-

vious subsection will be then discussed below in more detail. The ECB publishes a list of bonds purchased and held under the CSPP with the respective purchase dates. This list

includes each bond's International Securities Identification Number (ISIN) and is updated on a weekly basis. Similar to other asset purchase programmes, CSPP purchases are announced ex post which allows the exploitation of the official intervention as an exogenous, unexpected reduction in the supply of corporate bonds traded amongst investors.

As CDS premia are written on a single company, comprising a set of multiple bonds, I will have to consider the reference or parent company of each purchased bond in order to construct a reasonable benchmark for the DID estimation. Thus, the underlying data is retrieved from Bloomberg by matching each bond with its ultimate parent company using the ISINs. From Bloomberg I also collect information that includes bond level characteristics such as amount outstanding, coupons, country, currency, payment rank and maturity to redemption. Additionally, I obtain the rating at launch of the corporate bond issuances from four rating agencies S&P, Moody's, Fitch Ratings and Dominion Bond Rating Services. I follow *Grosse-Rueschkamp et al. (2017)* and use the credit rating at bond level as a proxy for the rating of the ultimate parent issuer. The implicit assumption is that the parent issuer rating is positively correlated with issue ratings, and further, that credit ratings are positively correlated across rating agencies. In total, there are 72 firms in my sample.

I use Markit as the central source for the CDS data. Markit offers comprehensive pricing data, collected directly from market makers, and subject to a rigorous data cleaning process.³⁷ However, matching the ECB's list of purchased bonds with the CDS data is not straightforward. The main concern is that in Markit CDS prices are not mapped to single ISINs but identified through a unique 6-digit REDCode number assigned by Markit for each reference company. As mentioned previously, CDS contracts are written on issuers and not appreciated at the issue level. Accordingly, prices are only available at the company level. Therefore, to avoid collecting data by hand, an intermediate step is required. I need to match the multiple bonds to the 6-digit REDCode with the aid of the ISINs. Only then am I able to match reference entities between the two data sources, whereby the matching procedure here will be based on the 6-digit REDCode.

After merging datasets and removing missing observations, I am able to identify 52 firms with available CDS data for the sample period between January 2015 and November 2017. The series covers the quoted spread, the reference company, the seniority tier,³⁸ the currency, the country, the industry, the recovery rate³⁹ and the restructuring clause levels of the respective CDS on a daily basis. In addition, 16

firms are dismissed because in the corresponding industries no bonds have been purchased by the ECB, or because the spread data does not cover the period from January 2015 until November 2017. In total, 36 firms remain in the sample.

Most noteworthy, there is no one-to-one correspondence between a CDS spread and its underlying entity, as the 6-digit REDCode identifies CDS contracts for all available restructuring and seniority characteristics. I follow *Berndt and Obreja (2010)* and *Mertens (2017)* who select 5-year contracts with modified-modified restructuring clause for senior unsecured Euro denominated debt. These are considered to be the most liquid CDS contracts in the European market.⁴⁰ As discussed earlier, restructuring events are less straightforward as compared to other credit events. Nevertheless, for my estimation I will focus on the modified-modified restructuring clause, which is most popular in Europe. It imposes a maturity limit of 60 months for restructured obligations and 30 months for all other obligations (*Berndt et al., 2007, Augustin et al., 2014*). My approach is motivated by recent evidence. In the absence of restructuring as a credit event, lenders to a reference company who also trade CDS linked to that same reference company - known as empty creditors - are likely to be tougher during debt renegotiations, refusing private workouts and making distressed borrowers more vulnerable to bankruptcy. For example, buyers of 'no restructuring' CDS contracts with bankruptcy as a credit event will only be paid if the reference firm files for bankruptcy (*Pollack, 2003, Subrahmanyam et al., 2014*). For a formal model see *Bolton & Oehmke (2011)*. One implication of the model is that the empty creditor problem is, in fact, priced in CDS premia. Hence, to avoid any distortion of results I will opt for restructuring as a credit event, particularly for the modified-modified restructuring clause. It should be noted that the overall tightening of the data comes at the expense of severely reducing the number of observations. As an example, consider the data for the company Aegon N.V. In the full dataset, there are 4009 observations for the 5-year maturity CDS, but after filtering for the seniority tier and the restructuring clause, only 743 remain.

After applying this filter, I am able to construct a representative sample of the treatment and control group. Hereby, within each industry, entities issuing peri-CSPP bonds serve as the treatment variable while the rest are assigned to the control group (see Figure 7). The key point, though, is to avoid the assignment to the control group occurring arbitrarily. Hence, for each treated firm I select only comparable firms from the same industry. More specifically, comparable firms are defined as those firms that - with respect to the CDS spread series - exhibit a similar pre-treatment trend as the treated group (see Appendix). The parallel trend assumption is, as already stated, an important prerequisite underlying the DID estimation. Overall, of the 36 firms, 9 can be assigned to the treatment group and 23 firms can be assigned to the con-

³⁷Note that I do not observe bid and ask quotes for CDS spreads, but only mid quotes. In particular, Markit reports spreads that are obtained by averaging the quotes reported by various financial institutions, inter-dealer brokers, and electronic trading platforms (*Giglio, 2014*). Moreover, reported CDS quotes reflect the sell-side offering price and not the finally agreed price between counterparties (*Liu et al., 2017*).

³⁸The credit risk a CDS references is not limited to a particular bond or loan, but comprises a predetermined set of debt obligations. Markit defines by means of the seniority tier the level of risk of these debt claims.

³⁹The recovery rate corresponds to the industry standard value of 40 percent for all CDS contracts in the sample (*Chen et al., 2010*).

⁴⁰*Blanco et al. (2005), Longstaff et al. (2005)* and *Norden and Weber (2009)* also choose the benchmark maturity of five years.

trol group. Table 2 presents characteristics of the 32 firms included in the final sample. The table lists the reference entities, together with basic descriptive information, such as the S&P credit rating, the country and the currency, as well as the daily spread average at the firm and industry level, and finally the number of observations in the CDS series. Each industry group is well diversified across ratings and country of risk, whereby the country selection is due to data availability. Control firms, on average, have slightly higher spreads as compared to treated firms (74.538 basis points versus 72.300 basis points).

It may be argued that the sample size of 32 firms in this context is not realistic. But according to a report by the International Capital Market Association the non-sovereign, non-financial CDS sector is in general modestly represented, while sovereign and financial CDS dominate the European CDS market. In fact, government and financial CDS contribute to 58 percent of the total notional outstanding as of 29 September 2017 (Callsen and Hill, 2018). Hence the data I use is not ideal but the best currently available for my purpose.

5. Empirical Results

This section reports the empirical findings regarding the CSPP impact on CDS prices. To recap, this paper seeks to examine two main questions.

- Do CDS rates increase relative to the time before the CSPP announcement?
- Do CDS rates increase less for CSPP-eligible firms relative to the control group (eligible firms not yet subject to the CSPP) and relative to the time before the CSPP announcement?

The regression specified in the previous section is estimated on daily CDS spreads using the DID approach in combination with robust standard errors to account for heteroscedasticity. The main regressor is the interaction term of the bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and the time dummy variable $Post_t$ (=1 from the initial purchase date and thereafter), that indicates the post-purchase period. The evaluation period is seven trading days before and after the purchase date of the respective bond. Empirical results are listed in Table 3 to 7, for each industry separately. Columns (1) to (3) show the results with $CDS\ Spread_{it}$ as the main dependent variable. Columns (4) to (6) show the results with $\log CDS\ Spread_{it}$ as the main dependent variable. This pattern is analogous for Table 3 to 7.

Column (1) starts with the specification that only includes the post-programme dummy term, $Post_t$, while controlling for unobserved time-invariant firm characteristics. The empirical result for the Basic Materials sector suggests that - in line with the first hypothesis - after the bond purchase ($Post_t$), all entities in this industry experience an average decrease of 9 basis points in their CDS spreads. In column

(2) the model is then further saturated with the interaction term $Treat_{it} Post_t$. Here the results are strengthened with a value of around 11 basis points for $Post_t$ (see Table 3). Similarly, for the Industrials sector CDS rates drop by up to more than 5 basis points (see Table 7). For the remaining three sectors CDS rates do not change as appreciably with the CSPP implementation. Corresponding figures are between 0 and 3 basis points (see Table 4 to 6). Hence, firms in the Basic Materials sector experience the largest fall in CDS rates, followed by the Industrial sector. In all instances, the estimates of the coefficient $Post_t$ are negative. In terms of statistical significance, daily CDS rates across industries are significantly more negative, after the CSPP purchase day than previously, in 7 out of 9 instances. These results hint towards a systematic decline in spreads following CSPP purchases.

In order to be able to infer causality I take a closer look at the interaction term $Treat_{it} Post_t$ in column (2). In agreement with the second conjecture, only a positive sign of the term $Treat_{it} Post_t$ implies that after the purchase date, and conditional on being purchased, spreads of treated entities drop less than spreads of their non-CSPP counterparts. For the Basic Materials sector this difference in drop is indeed positive and robust to the inclusion of entity fixed effects. Particularly, CDS rates have decreased post-CSPP purchase by almost 9 basis points less for the treated firm relative to control group firms and relative to the pre-CSPP event. Hence, this finding indicates the existence of spillover effects of the CSPP programme.

However, it is possible that time-specific shocks are driving the results. Column (3) controls for time-specific shocks (time fixed effects) as defined in Equation (2). Naturally, the post-programme dummy term $Post_t$ is dropped when adding day fixed effects, as the programme affects all entities at a specific point in time. The magnitude of the coefficient still remains fairly constant across the specification and the standard error does not vary significantly. Thus, for the Basic Material sector, spillover effects are indeed evident (see Table 3). However, across industries there is no clear pattern with respect to the sign of the coefficient of the interaction term, as detected for $Post_t$. A similar effect can only be observed for one subset of the Industrial sector. In fact, the bond purchase on April 30, 2017 prompts CDS rates of the treatment group (Atlantia S.p.A.) to decline by 4 basis points less as compared to the control group (see Table 7). For the remaining three industries corresponding coefficients are either statistically insignificant or the magnitude is negligible in economic terms (below 1 basis point). Note that in 6 out of 9 cases where the interaction term is statistically significant, 3 carry a positive sign and 3 a negative sign. Two important implications can be drawn from these findings. First, the for the most part low but highly significant estimates implicate that there is actually an association between the CSPP and CDS rates. Second, irregularities with respect to the sign of the coefficients do not allow for interpretations on spillover effects as postulated in the second hypothesis.

However, it might be the case that the prior relationship is non-linear. To account for that I run the same regression but

Table 2: Sample Description

The table summarizes the final database after filtering, comprising 5-year CDS denominated in EUR for France, Germany, Greece, Italy, Netherlands, Portugal, Spain and the United Kingdom in a period ranging from January 2015 to November 2017. Mean spreads are calculated in basis points; (1) purchased on August 8th, 2016 (2) purchased on August 15th, 2016 (3) purchased on October 3rd, 2016 (4) purchased on January 23rd, 2017 (5) purchased on May 1st, 2017.

Entity	Currency	Country	Rating S&P	Observations	Mean
Panel A: Basic Materials					
Koninklijke DSM N.V.	EUR	Netherlands	A-	743	40.142
LINDE Aktiengesellschaft	EUR	Germany	BBB	743	30.103
LANXESS Aktiengesellschaft (1)	EUR	Germany	BBB	743	71.885
XSTRATA LIMITED	EUR	United Kingdom	BBB+	743	244.212
Panel B: Financials					
Aegon N. V. (1)	EUR	Netherlands	A-	743	87.929
Allianz SE	EUR	Germany	AA-	743	38.185
ASSICURAZIONI GENERALI	EUR	Italy	A+	743	110.186
AXA	EUR	France	A-	743	58.702
NN Group N. V.	EUR	Netherlands	BBB	743	70.018
UNIBAIL-RODAMCO SE	EUR	France	A	743	60.592
Panel C: Industrials					
Airbus Group SE	EUR	Netherlands	BBB+	479	65.538
Airbus Group N.V.	EUR	Netherlands	BBB+	110	51.858
ATLANTIA S. P. A. (1),(5)	EUR	Italy	A	743	66.599
BRISA - AUTO-ESTRADAS DE PORTUGAL, S. A.	EUR	Portugal	A	742	126.100
BRISA - CONCESSAO RODOVIARIA, S. A.	EUR	Portugal	A	743	125.810
HeidelbergCement AG (4)	EUR	Germany	BBB-	743	113.051
THALES	EUR	France	BBB+	743	51.187
Lafarge	EUR	France	BB+	743	54.775
PostNL N. V. (2)	EUR	Netherlands	BBB+	743	54.281
Siemens Aktiengesellschaft	EUR	Germany	AA-	743	36.376
COMPAGNIE DE SAINT-GOBAIN	EUR	France	BBB+	743	63.776
VINCI	EUR	France	BBB+	743	52.595
Panel D: Telecommunications Services					
Deutsche Telekom AG	EUR	Germany	BBB+	743	45.348
Orange (1)	EUR	France	A-	743	58.990
Vivendi	EUR	France	BBB	743	62.271
Panel E: Utilities					
EnBW Energie Baden-Wuerttemberg AG	EUR	Germany	A-	743	53.941
ENEL S. P. A.	EUR	Italy	A-	743	86.083
ENGIE (1)	EUR	France	A-	589	55.754
E.ON SE (1)	EUR	Germany	AA-	743	73.969
EDISON S. P. A. (3)	EUR	Italy	BBB	743	64.815
Iberdrola S.A.	EUR	Spain	BBB+	743	72.735
RWE Aktiengesellschaft	EUR	Germany	A+	708	92.226
					71.360

Table 3: CSPP Effects on CDS Spreads in Basic Materials Sector (7 Trading Days)

Table 3 presents the DID regression for the subsample of entities within the Basic Materials sector. This table checks whether entities within the treated sample are affected differently in terms of CDS spreads relative to the control group. The dependent variable is the corporate CDS spread. The main regressor is an interaction term of a bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and a time dummy variable $Post_t$, that indicates the purchase of the respective bond under the CSPP (=1 from August 8, 2016 and after). Columns (1) to (3) correspond to the linear regression of Eq. 2. Columns (3) to (6) use the semi log regression of Eq. 3. Observations are between 7 trading days before and after August 8, 2016 and the bond purchased has been issued by Lanxess AG; Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Variables	$CDS\ Spread_{it}$			$\log(CDS\ Spread_{it})$		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-8.934*** (2.019)	-11.09*** (2.625)		-0.0680*** (0.00782)	-0.0772*** (0.00995)	
$Treat_{it}Post_t$		8.615*** (2.643)	8.615*** (3.074)		0.0370*** (0.0111)	0.0370*** (0.0106)
Observations	56	56	56	56	56	56
Entities	4	4	4	4	4	4
R-squared	0.994	0.994	0.995	0.999	0.999	0.999
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES

Table 4: CSPP Effects on CDS Spreads in Telecommunications Services Sector (7 Trading Days)

Table 4 presents the DID regression for the subsample of entities within the Telecommunications Services sector. This table checks whether entities within the treated sample are affected differently in terms of CDS spreads relative to the control group. The dependent variable is the corporate CDS spread. The main regressor is an interaction term of a bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and a time dummy variable $Post_t$, that indicates the purchase of the respective bond under the CSPP (=1 from August 8, 2016 and after). Columns (1) to (3) correspond to the linear regression of Eq. 2. Columns (3) to (6) use the semi log regression of Eq. 3. Observations are between 7 trading days before and after August 8, 2016 and the bond purchased has been issued by Orange S.A.; Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Variables	$CDS\ Spread_{it}$			$\log(CDS\ Spread_{it})$		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-1.471*** (0.205)	-1.350*** (0.243)		-0.0276*** (0.00389)	-0.0257*** (0.00473)	
$Treat_{it}Post_t$		-0.364 (0.453)	-0.364** (0.140)		-0.00551 (0.00843)	-0.00551** (0.00221)
Observations	42	42	42	42	42	42
Entities	3	3	3	3	3	3
R-squared	0.995	0.995	1.000	0.995	0.995	1.000
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES

on the logarithm of the dependent variable ($\log CDS\ Spread_{it}$) as defined in Equation (3). The initial results are broadly unchanged. For the most part, the $Post_t$ coefficients in column (4) and (5) keep their signs and remain highly significant across industries, when previously designated as statistically significant. Again, the effects are most striking for the Basic Materials sector as presented in Table 3. Following the bond purchase by the ECB, within this industry, CDS spreads exhibit on average a decrease between 7 and 8 percent. Most noteworthy, the CDS spread for the treatment entity decreases by 3.77 percentage points less relative to the control group after the announcement of the CSPP. Likewise, within the Industrials sector (for the issuer Atlantia S.p.A.) CDS spreads drop on average by 6 percent post-CSPP. This

finding is highly significant and holds for both, the purchase date in 2016 and 2017. Interestingly, only for the latter purchase date the coefficient on the interaction term is positive and sizeable with a figure of roughly 6 percentage points (see Table 7).

For the CDS rates of issuers Engie S.A. and E.ON International Finance B.V. from the Utilities sector there is a significant decline of the order of around 4 percent after the ECB conducts the bond purchases, whereas post-CSPP and relative to the control group the effect between the treated and control group becomes statistically indistinguishable from zero (see Table 6). The Financials sector experiences a drop in spreads of similar magnitude, but again the interaction is not statistically different from zero (see Table 5). Moreover,

Table 5: CSPP Effects on CDS Spreads in Financials Sector (7 Trading Days)

Table 5 presents the DID regression for the subsample of entities within the Financials sector. This table checks whether entities within the treated sample are affected differently in terms of CDS spreads relative to the control group. The dependent variable is the corporate CDS spread. The main regressor is an interaction term of a bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and a time dummy variable $Post_t$, that indicates the purchase of the respective bond under the CSPP (=1 from August 8, 2016 and after). Columns (1) to (3) correspond to the linear regression of Eq. 2. Columns (3) to (6) use the semi log regression of Eq. 3. Observations are between 7 trading days before and after August 8, 2016 and the bond purchased has been issued by Aegon N.V.; Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Variables	$CDS\ Spread_{it}$			$\log(CDS\ Spread_{it})$		
	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-3.147*** (0.364)	-3.118*** (0.416)		-0.0430*** (0.00455)	-0.0447*** (0.00528)	
$Treat_{it}Post_t$		-0.172 (0.818)	-0.172 (0.584)		0.0104 (0.00902)	0.0104 (0.00677)
Observations	84	84	84	84	84	84
Entities	6	6	6	6	6	6
R-squared	0.997	0.997	0.998	0.997	0.997	0.998
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES

Table 6: CSPP Effects on CDS Spreads in Utilities Sector (7 Trading Days)

Table 6 presents the DID regression for the subsample of entities within the Utilities sector. This table checks whether entities within the treated sample are affected differently in terms of CDS spreads relative to the control group. The dependent variable is the corporate CDS spread. The main regressor is an interaction term of a bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and a time dummy variable $Post_t$ that indicates the purchase of the respective bond under the CSPP. Columns (1) to (3) correspond to the linear regression of Eq. 2. Columns (3) to (6) use the semi log regression of Eq. 3. Observations are between 7 trading days before and after August 8, 2016 and October 3, 2016, for Panel A and B respectively. Purchased bonds are issued by Engie S.A. and E.ON International Finance B.V. as shown in Panel A; and by Edison S.p.A. shown in Panel B; Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: 8 August 2016		$CDSSpread_{it}$			$\log(CDSSpread_{it})$		
Variables	(1)	(2)	(3)	(4)	(5)	(6)	
$Post_t$	-2.603*** (0.188)	-2.740*** (0.232)		-0.0396*** (0.00261)	-0.0390*** (0.00311)		
$Treat_{it}Post_t$		0.480 (0.385)	0.480** (0.224)		-0.00222 (0.00576)	-0.00222 (0.00215)	
Observations	98	98	98	98	98	98	
Entities	7	7	7	7	7	7	
R-squared	0.996	0.996	0.999	0.997	0.997	1.000	
Entity FE	YES	YES	YES	YES	YES	YES	
Time FE			YES			YES	
Panel B: 3 October 2016		$CDSSpread_{it}$			$\log(CDSSpread_{it})$		
Variables	(1)	(2)	(3)	(4)	(5)	(6)	
$Post_t$	0.447** (0.221)	0.622** (0.272)		0.00673** (0.00290)	0.00918** (0.00356)		
$Treat_{it}Post_t$		-0.877*** (0.292)	-0.877** (0.354)		-0.0123*** (0.00378)	-0.0123** (0.00481)	
Observations	70	70	70	70	70	70	
Entities	5	5	5	5	5	5	
R-squared	0.996	0.996	0.998	0.997	0.997	0.999	
Entity FE	YES	YES	YES	YES	YES	YES	
Time FE			YES			YES	

Table 7: CSPP Effects on CDS Spreads in Industrials Sector (7 Trading Days)

Table 7 presents the DID regression for the subsample of entities within the Industrials sector. This table checks whether entities within the treated sample are affected differently in terms of CDS spreads relative to the control group. The dependent variable is the corporate CDS spread. The main regressor is an interaction term of a bond CSPP purchase dummy $Treat_{it}$ (=1 if bond of entity i is purchased at day t) and a time dummy variable $Post_t$, that indicates the purchase of the respective bond under the CSPP. Columns (1) to (3) correspond to the linear regression of Eq. 2. Columns (3) to (6) use the semi log regression of Eq. 3. Observations are between 7 trading days before and after the specific purchase date as shown in Panel A to D respectively. Purchased bonds are issued by Atlantia S.p.A. as shown in Panel A and B; and by PostNL N.V. and HeidelbergCement Finance B.V. shown in Panel C and D respectively; Robust standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Panel A: 8 August 2016		$CDSSpread_{it}$			$log(CDSSpread_{it})$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-5.272*** (0.468)	-5.343*** (0.525)		-0.0575*** (0.00344)	-0.0558*** (0.00375)	
$Treat_{it}Post_t$		0.639 (0.707)	0.639 (0.586)		-0.0151* (0.00802)	-0.0151*** (0.00435)
Observations	126	126	126	126	126	126
Entities	9	9	9	9	9	9
R-squared	0.997	0.997	0.997	0.999	0.999	0.999
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES
Panel B: 1 May 2017		$CDSSpread_{it}$			$log(CDSSpread_{it})$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-3.889*** (0.467)	-4.292*** (0.499)		-0.0577*** (0.00625)	-0.0638*** (0.00667)	
$Treat_{it}Post_t$		4.030*** (0.929)	4.030*** (0.699)		0.0611*** (0.0109)	0.0611*** (0.0105)
Observations	140	140	140	140	140	140
Entities	10	10	10	10	10	10
R-squared	0.994	0.994	0.997	0.994	0.995	0.998
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES
Panel C: 15 August 2016		$CDSSpread_{it}$			$log(CDSSpread_{it})$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	-1.591*** (0.399)	-1.624*** (0.457)		-0.0202*** (0.00348)	-0.0190*** (0.00391)	
$Treat_{it}Post_t$		0.268 (0.515)	0.268 (0.626)		-0.00984 (0.00632)	-0.00984** (0.00459)
Observations	112	112	112	112	112	112
Entities	8	8	8	8	8	8
R-squared	0.998	0.998	0.998	0.999	0.999	0.999
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES
Panel D: 23 Januar 2017		$CDSSpread_{it}$			$log(CDSSpread_{it})$	
Variables	(1)	(2)	(3)	(4)	(5)	(6)
$Post_t$	0.0719 (0.323)	0.176 (0.356)		0.00459 (0.00427)	0.00628 (0.00471)	
$Treat_{it}Post_t$		-0.939 (0.722)	-0.939** (0.452)		-0.0151* (0.00869)	-0.0151** (0.00611)
Observations	126	126	126	126	126	126
Entities	9	9	9	9	9	9
R-squared	0.998	0.998	0.999	0.997	0.998	0.998
Entity FE	YES	YES	YES	YES	YES	YES
Time FE			YES			YES

following the CSPP policy, for Telecommunication companies the sign on the interaction estimate is reversed and negative, which seems to negate the existence of spillover effects in connection with the second assumption. However, the figure is relatively small in economic terms (see Table 4). Overall, the most pronounced impact in lowering credit risk can be observed for the sector of Basic Materials with a decrease of 8 percent. Noticeable evidence on spillover effects can be inferred from the Industrial sector with a value equal to 6 percentage points. In 7 out of 9 instances where the interaction term is statistically significant, 2 carry a positive sign and 5 carry a negative sign. Despite assuming non-linearity, the results of this exercise do not contribute to further insights. These findings rather suggest that spillover effects are limited to specific bonds (which is most evident for the issuer Atlantia S.p.A.). Note that the large R-squared values throughout all specifications are based on the fact that fixed effects often capture a lot of the variation in the data.

6. Discussion

The CSPP was designed to complement the main thrust of ECB'S QE policy. The overall goal has been to ease financial conditions for corporates, and ultimately to support a sustained economic recovery in the euro area. This paper adds to the strand of literature to study the CSPP impact, and especially the spillover effects of monetary policy decisions on related financial markets. While it is difficult to be certain about the effects of the CSPP policy without a greater body of experience than is so far available, some provisional conclusions may be possible.

To summarize, I find that, consistent with the initial assumption, the CSPP programme has contained credit risk across European non-financial corporates. The results indicate that credit market reactions to the CSPP event - measured by means of CDS prices - imply negative CDS rates throughout. In contrast, spillover effects to non-CSPP firms have been heterogeneous within and across industries. The empirical support for the second conjecture is limited, and if anything, rather bond specific. Hence, the ECB'S commitment to continue the CSPP is indeed helping to lift credit constraints overall, but according to my estimation the programme seems to not have stronger effects on firms not subject to the CSPP, as suggested by previous work.

A potential critique of the above analysis is related to the identification strategy. While the second hypothesis is motivated by the fact that the reduction in CDS rates - prompted by the CSPP - spills over to riskier CDS instruments, in the estimation the full sample is restricted to the CSPP portfolio. In other words, the implicit assumption underlying the empirical strategy is that control group firms transferred later to the CSPP portfolio are higher credit risk firms and thus non-CSPP-eligible which is, however, not necessarily true. As elaborated earlier, the empirical strategy is convenient as it dismisses any endogeneity concerns. Nonetheless, defining a too narrow control group may lead to inconsistent results.

In this respect, the sample could be extended to select Eurozone investment grade-rated companies issuing USD denominated bonds into the control group. An alternative control group may comprise European investment grade-rated firms that are incorporated in countries outside of the Eurozone. These approaches would still alleviate any endogeneity concerns, as treated companies would most likely not differ systematically from the control group except by the currency or country. Next to that, as in my estimation there are only 32 firms with available CDS spreads, an adjustment of the sample as proposed would indeed allow access to a larger database which may in turn complement the prior results. However, such analyses are beyond the scope of this paper, leaving space for future research.

At the aggregate level, it may well be argued that CDS contracts as such are restrictive for the evaluation of the CSPP impact. In fact, CDS do not refer to a single bond but to a firm, which issues various bonds. And indeed, the prevailing examination hints towards the fact that CSPP purchases are having a systematic effect at the individual bond-level, whereas the effect on the entity as a whole is rather ambiguous. Beyond that there are some other shortcomings associated with this financial instrument. As discussed earlier, CDS rates present rather an upper limit on the price of credit risk. In fact, [Subrahmanyam et al. \(2014\)](#) find that the inception of CDS increases the credit risk of underlying reference entities due to the higher likelihood of credit rating downgrades and bankruptcy. The increase in credit risk is also associated with the absence of borrower monitoring and tougher debt renegotiations. Similarly, [Arce et al. \(2017\)](#) report that the cost of debt of risky firms actually increases after CDS trading is initiated. Hence, the choice of CDS as the variable of interest may lead to distorted results because estimation results, if anything, will be biased upwards.

However, this line of reasoning omits a certain important aspect, that is that the CSPP announcement date absorbs pricing-relevant information for the most part, generating lower bound estimates for the individual purchase date ([Arce et al., 2017](#)). In aggregate, there exists the issue of an overestimation on the one hand, and an underestimation on the other hand. Prospective research will be required to disentangle these two effects, and disclose whether they offset each other.

Within this context, it is also crucial to understand that the predominant focus on CDS quotes may be too simplistic. Instead, a shift towards a composite dataset of both CDS transaction data and CDS quotes may provide a more comprehensive picture of CDS activities, revealing supplementary information on referenced firms. Further, as a robustness check future research may consider multi-name CDS instruments which typically represent the more liquid part of the relevant single-name CDS market ([Fontana and Scheicher, 2016](#)). Corresponding CDS spreads would then serve as a more powerful indicator of credit risk.

References

- Abad, J., Aldasoro, I., Aymanns, C., D'Errico, M., Rousová, L., Hoffmann, P., Langfield, S., Neychev, M., and Roukny, T. Shedding light on dark markets: First insights from the new EU-wide OTC derivatives dataset. Technical report, European Systemic Risk Board Occasional Paper Series No. 11, 2016.
- Abidi, N., Flores, I. M., and Eterovic, N. A. Who benefits from the ECB's Corporate Sector Purchase Programme? a difference-in-discontinuities approach. Technical report, Working Paper, 2017.
- Acharya, V. V. and Johnson, T. C. Insider trading in credit derivatives. *Journal of Financial Economics*, 84(1):110–141, 2007.
- Altavilla, C., Carboni, G., and Motto, R. Asset purchase programmes and financial markets: lessons from the euro area. *ECB Working Paper Series No. 1864*, 2015.
- Andrade, P., Breckenfelder, J. H., De Fiore, F., Karadi, P., and Tristani, O. The ECB's asset purchase programme: an early assessment. *ECB Working Paper Series No. 1956*, 2016.
- Arce, O., Gimeno, R., and Mayordomo, S. Making room for the needy: The credit-reallocation effects of the ECB's Corporate QE. *Banco de España Working Paper No. 1743*, 2017.
- Ashenfelter, O. Estimating the effect of training programs on earnings. *The Review of Economics and Statistics*, 6(1):47–57, 1978.
- Augustin, P., Subrahmanyam, M. G., Tang, D. Y., and Wang, S. Q. Credit default swaps: A survey. *Foundations and Trends in Finance*, 9(1–2):1–196, 2014.
- Batta, G. E., Qiu, J., and Yu, F. Credit derivatives and analyst behavior. *The Accounting Review*, 91(5):1315–1343, 2016.
- Beau, E., Hill, J., Hussain, T., and Nixon, D. Bank funding costs: what are they, what determines them and why do they matter? *Bank of England Quarterly Bulletin*, 54(4):370–384, 2014.
- Beber, A., Brandt, M. W., and Kavajecz, K. A. Flight-to-quality or flight-to-liquidity? evidence from the euro-area bond market. *The Review of Financial Studies*, 22(3):925–957, 2009.
- Beirne, J., Dalitz, L., Ejsing, J., Grothe, M., Manganelli, S., Monar, F., Sahel, B., Susec, M., Tapking, J., and Vong, T. The impact of the eurosystem's covered bond purchase programme on the primary and secondary markets. *ECB Occasional Paper Series*, (122), 2011.
- Bernanke, B. S. The economic outlook and monetary policy. Speech at the Federal Reserve Bank of Kansas City Economic Symposium Jackson Hole, Wyoming, 27 August 2010, 2010.
- Bernanke, B. S. and Reinhart, V. R. Conducting monetary policy at very low short-term interest rates. *American Economic Review*, 94(2):85–90, 2004.
- Berndt, A. and Obreja, I. Decomposing european CDS returns. *Review of Finance*, 14(2):189–233, 2010.
- Berndt, A. and Ostrovnaya, A. Do equity markets favor credit market news over options market news? *The Quarterly Journal of Finance*, 4(02), 2014.
- Berndt, A., Jarrow, R. A., and Kang, C.-O. Restructuring risk in credit default swaps: An empirical analysis. *Stochastic Processes and their Applications*, 117(11):1724–1749, 2007.
- Bertrand, M., Duflo, E., and Mullainathan, S. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275, 2004.
- Bini Smaghi, L. Conventional and unconventional monetary policy. Keynote speech at the International Center for Monetary and Banking Studies, Geneva, 28 April 2009, 2009.
- Blanco, R., Brennan, S., and Marsh, I. W. An empirical analysis of the dynamic relation between investment-grade bonds and credit default swaps. *The Journal of Finance*, 60(5):2255–2281, 2005.
- Breitenfellner, B. and Wagner, N. Explaining aggregate credit default swap spreads. *International Review of Financial Analysis*, 22:18–29, 2012.
- Callsen, G. and Hill, A. The european corporate single name credit default swap market. *International Capital Market Association Working Paper*, 2018.
- Chen, R.-R., Fabozzi, F. J., and Sverdlow, R. Corporate credit default swap liquidity and its implications for corporate bond spreads. *Journal of Fixed Income*, 20(2):31–57, 2010.
- Coeure, B. Bond scarcity and the ECB's asset purchase programme. Speech at the Club de Gestion Financiere d'Associés en Finance, Paris, 3 April 2017, 2017.
- De Santis, R. A. Credit spreads, economic activity and fragmentation. *ECB Working Paper Series*, (1930), 2016a.
- De Santis, R. A. Impact of the asset purchase programme on euro area government bond yields using Bloomberg news. *ECB Working Paper Series*, (1939), 2016b.
- Delatte, A.-L., Gex, M., and López-Villavicencio, A. Has the CDS market influenced the borrowing cost of european countries during the sovereign crisis? *Journal of International Money and Finance*, 31(3):481–497, 2012.
- Demertzis, M. and Wolff, G. B. What impact does the ecb's quantitative easing policy have on bank profitability? Technical report, Bruegel Policy Contribution, 2016.
- Diamond, D. W. and Verrecchia, R. E. Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics*, 18(2):277–311, 1987.
- Draghi, M. Speech at the global investment conference. London, 26 July 2012, 2012.
- Draghi, M. Introductory statement at the hearing of the committee on economic and monetary affairs of the european parliament. Brussels, 26 September 2016, 2016.
- Duffie, D. Credit swap valuation. *Financial Analysts Journal*, 55(1):73–87, 1999.
- D'Amico, S. and King, T. B. Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics*, 108(2):425–448, 2013.
- Eggertsson, G. B. and Woodford, M. Zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 34(1):139–233, 2003.
- Eser, F. and Schwaab, B. Evaluating the impact of unconventional monetary policy measures: Empirical evidence from the ECB's securities markets programme. *Journal of Financial Economics*, 119(1):147–167, 2016.
- European Central Bank. The monetary policy of the ecb. 3rd Edition, May 2011, ECB, 2011.
- European Central Bank. The governing council's expanded asset purchase programme. Economic Bulletin Box 1 (1), 2015.
- European Central Bank. The corporate bond market and the ECB's corporate sector purchase programme. Economic Bulletin Box 2 (5), 2016.
- European Central Bank. The ECB's corporate sector purchase programme: its implementation and impact. Economic Bulletin Box 2 (4), 2017a.
- European Central Bank. The recalibration of the ECB's asset purchase programme. Economic Bulletin Box 2 (7), 2017b.
- European Central Bank Website. Corporate sector purchase programme (CSPP) - Questions & answers, 2018a. URL <https://www.ecb.europa.eu/mopo/implement/omt/html/cspp-qa.en.html>. last accessed: 3 June 2018.
- European Central Bank Website. Monetary policy - introduction, 2018b. URL <https://www.ecb.europa.eu/mopo/intro/html/index.en.html>. last accessed: 3 June 2018.
- Fama, E. F. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.
- Fawley, B. W. and Neely, C. J. Four stories of quantitative easing. *Federal Reserve Bank of St. Louis Review*, 95(1):51–88, 2013.
- Fontana, A. and Scheicher, M. An analysis of euro area sovereign CDS and their relation with government bonds. *ECB Working Paper Series*, (1271), 2016.
- Forte, S. and Pena, J. I. Credit spreads: An empirical analysis on the informational content of stocks, bonds, and CDS. *Journal of Banking & Finance*, 33(11):2013–2025, 2009.
- Fratzscher, M., Duca, M. L., and Straub, R. ECB unconventional monetary policy: Market impact and international spillovers. *IMF Economic Review*, 64(1):36–74, 2016.
- Fukunaga, I., Kato, N., and Koeda, J. Maturity structure and supply factors in japanese government bond markets. *Institute For Monetary And Economic Studies Discussion Paper*, (2015-E-10), 2015.
- Fulop, A. and Lescourret, L. How liquid is the CDS market? In *Working Paper*. Citeseer, 2007.
- Gagnon, J., Raskin, M., Remache, J., and Sack, B. The financial market effects of the federal reserve's large-scale asset purchases. *International Journal of Central Banking*, 7(1):3–43, 2011.
- Giannone, D., Lenza, M., Pill, H., and Reichlin, L. Non-standard monetary policy measures and monetary developments. *ECB Working Paper Series*, (1290), 2011.
- Giglio, S. Credit default swap spreads and systemic financial risk. *Havard University Working Paper*, 10(9):104–141, 2014.

- Gilchrist, S. and Zakrajšek, E. The impact of the federal reserve's large-scale asset purchase programs on corporate credit risk. *Journal of Money, Credit and Banking*, 45(s2):29–57, 2013.
- Gormley, T. A. and Matsa, D. A. Common errors: How to (and not to) control for unobserved heterogeneity. *The Review of Financial Studies*, 27(2):617–661, 2013.
- Grosse-Rueschkamp, B., Steffen, S., and Streitz, D. Cutting out the middleman: The ECB as Corporate Bond Investor. Available at SSRN, 2017.
- Hamilton, J. D. and Wu, J. C. The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking*, 44(1):3–46, 2012.
- Hancock, D. and Passmore, W. Did the federal reserve's MBS purchase program lower mortgage rates? *Journal of Monetary Economics*, 58(5):498–514, 2011.
- Hilscher, J., Pollet, J. M., and Wilson, M. Are credit default swaps a sideshow? evidence that information flows from equity to CDS markets. *Journal of Financial and Quantitative Analysis*, 50(3):543–567, 2015.
- Hull, J. C. and White, A. Valuing credit default swaps I: No counterparty default risk. *Journal of Derivatives*, 8(1):29–40, 2000.
- Jankowitsch, R., Pullirsch, R., and Veža, T. The delivery option in credit default swaps. *Journal of Banking & Finance*, 32(7):1269–1285, 2008.
- Jarrow, R. A. Problems with using CDS to infer default probabilities. *Journal of Fixed Income*, 21(4):6–12, 2012.
- Joyce, M., Lasaosa, A., Stevens, I., and Tong, M. The financial market impact of quantitative easing in the United Kingdom. *International Journal of Central Banking*, 7(3):113–161, 2011a.
- Joyce, M., Tong, M., and Woods, R. The United Kingdom's quantitative easing policy: design, operation and impact. *Bank of England Quarterly Bulletin*, 51(3):200–212, 2011b.
- Joyce, M., Miles, D., Scott, A., and Vayanos, D. Quantitative easing and unconventional monetary policy—an introduction. *The Economic Journal*, 122(564):F271–F288, 2012.
- Joyce, M. A. and Tong, M. Qe and the gilt market: a disaggregated analysis. *The Economic Journal*, 122(564):F348–F384, 2012.
- Keynes, J. M. *The general theory of employment, interest, and money*. Harcourt Brace Jovanovich, London 1964 (reprint of the 1936 edition), 1936.
- Koijen, R. S., Koulischer, F., Nguyen, B., and Yogo, M. Quantitative easing in the euro area: The dynamics of risk exposures and the impact on asset prices. *Working Paper*, (601), 2016.
- Krishnamurthy, A. and Vissing-Jorgensen, A. The effects of quantitative easing on interest rates: Channels and implications for policy. Technical report, National Bureau of Economic Research, 2011.
- Krishnamurthy, A., Nagel, S., and Vissing-Jorgensen, A. ECB policies involving government bond purchases: Impact and channels. *NBER Working Paper*, (23985), 2017.
- Lam, W. R. Bank of Japan's monetary easing measures: Are they powerful and comprehensive? *IMF working papers*, (11/264):1–18, 2011.
- Lechner, M. The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics*, 4(3):165–224, 2011.
- Liu, J., Ng, J., Tang, D. Y., and Zhong, R. CDS trading and stock price crash risk. *Working Paper*, 2017.
- Longstaff, F. A., Mithal, S., and Neis, E. Corporate yield spreads: Default risk or liquidity? new evidence from the credit default swap market. *The Journal of Finance*, 60(5):2213–2253, 2005.
- Markit Group. Markit credit indices. a primer, 2008.
- McLaren, N., Banerjee, R. N., and Latto, D. Using changes in auction maturity sectors to help identify the impact of QE on gilt yields. *The Economic Journal*, 124(576):453–479, 2014.
- Meier, A. *Panacea, curse, or nonevent? Unconventional monetary policy in the United Kingdom*. Number 9-163. International Monetary Fund Working Paper, 2009.
- Mertens, R. The reactions of equity and credit markets to unconventional monetary policy—are the markets buying it? *Working Paper*, 2017.
- Merton, R. C. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470, 1974.
- Morris, S., Tödtling-Schönhofer, H., and Wiseman, M. Design And Commissioning Of Counterfactual Impact Evaluations - A Practical Guidance For ESF Managing Authorities. *Publications Office of the European Union, Luxembourg*, 2013.
- Neely, C. J. *The large scale asset purchases had large international effects*. Number Working Paper Series No. 2010-018C. Federal Reserve Bank of St. Louis, Research Division, 2012.
- Norden, L. and Weber, M. The co-movement of credit default swap, bond and stock markets: An empirical analysis. *European Financial Management*, 15(3):529–562, 2009.
- O'Kane, D., Pedersen, C. M., and Turnbull, S. M. *The restructuring clause in credit default swap contracts*, volume 2003-Q1/Q2. Lehman Brothers Quantitative Credit Research, 2003.
- Opsahl, T. and Newton, W. Credit risk and companies' inter-organizational networks: Assessing impact of suppliers and buyers on CDS spreads. *Working Paper*, 2015.
- Pollack, E. R. Assessing the usage and effect of credit derivatives. In *Harvard Law School International Finance Seminar (April)*, Working Paper, 2003.
- Regulation, E. P. Regulation (EU) 236/2012 of the European Parliament and of the Council of 14 March 2012 on short selling and certain aspects of credit default swaps. *Official Journal of the European Union L 86/1*, 2012.
- Schlepper, K., Hofer, H., Riordan, R., and Schrimpf, A. Scarcity effects of QE: A transaction-level analysis in the bund market. *BIS Working Papers*, (625), 2017.
- Skinner, F. S. and Townend, T. G. An empirical analysis of credit default swaps. *International Review of Financial Analysis*, 11(3):297–309, 2002.
- Stulz, R. M. Credit default swaps and the credit crisis. *Journal of Economic Perspectives*, 24(1):73–92, 2010.
- Subrahmanyam, M. G., Tang, D. Y., and Wang, S. Q. Does the tail wag the dog? the effect of credit default swaps on credit risk. *The Review of Financial Studies*, 27(10):2927–2960, 2014.
- Swanson, E. T. Let's twist again: a high-frequency event-study analysis of operation twist and its implications for QE2. *Brookings Papers on Economic Activity*, 42(1):151–188, 2011.
- Szczerbowicz, U. et al. The ECB unconventional monetary policies: have they lowered market borrowing costs for banks and governments? *International Journal of Central Banking*, 11(4):91–127, 2015.
- Tamura, K. and Tabakis, E. The use of credit claims as collateral for eurosystem credit operations. *ECB Occasional Paper Series*, (148), 2013.
- Tang, D. Y. and Yan, H. Liquidity and credit default swap spreads. *Working Paper*, 2007.
- Taylor, J. B. Discretion versus policy rules in practice. In *Carnegie-Rochester Conference Series on Public Policy*, volume 39, pages 195–214. Elsevier, 1993.
- Trichet, J.-C. The ECB's enhanced credit support. Keynote address at the University of Munich, Munich, 13 July 2009, 2009.
- Ueda, K. The effectiveness of non-traditional monetary policy measures: The case of the Bank of Japan. *The Japanese Economic Review*, 63(1):1–22, 2012.
- Vayanos, D. and Vila, J.-L. A preferred-habitat model of the term structure of interest rates. Technical report, National Bureau of Economic Research, 2009.
- Woodford, M. Methods of policy accommodation at the interest-rate lower bound. *Speech at the Federal Reserve Bank of Kansas City Symposium Jackson Hole, Wyoming, 31 August 2012*, 2012.
- Wright, J. H. What does monetary policy do to long-term interest rates at the zero lower bound? *The Economic Journal*, 122(564):F447–F466, 2012.
- Zhu, H. An empirical comparison of credit spreads between the bond market and the credit default swap market. *Journal of Financial Services Research*, 29(3):211–235, 2006.