



Variance Risk Premia

Alexander Wahl

Ludwig-Maximilians-Universität München

Abstract

Using a relatively model-free approach to extract the risk-neutral expected variance from an extensive set of traded options on 29 single stocks and eight stock indices, I derive the variance risk premium defined as the difference between the actually realized variance and the expected variance under the risk-neutral measure. The analysis reveals that variance risk premia are persistently negative for the majority of underlyings and show a clear link to the underlying's exposure to systematic market variance. Moreover, I find that both the risk associated with continuous as well as discontinuous price movements contribute to observed variance risk premia.

Keywords: Variance risk premium, Volatility premium, Jumps, Risk-neutral, Model-free

1. Introduction

1.1. Problem outline

Even though the option pricing model invented by [Black and Scholes \(1973\)](#) is probably the most famous of its kind and often used in practice (albeit with modifications), it is well known that particularly the assumption of a constant return variance is not fulfilled in practice and volatility fluctuates over time. By now a plethora of studies (e.g. [Bakshi and Kapadia \(2003\)](#), [Coval and Shumway \(2001\)](#), [Carr and Wu \(2009\)](#)) document evidence that certain option positions earn average returns significantly different from zero although these positions are constructed to be insensitive to fluctuations in the price of the underlying. Indeed, this return pattern is most often interpreted as evidence that investors price the risk associated with stochastic volatility, which gives rise to the so-called variance risk premium.

Over the years, a veritable market for volatility products has emerged and different parties such as for example hedge funds (e.g. [Bondarenko \(2004\)](#)) engage in volatility selling strategies or are otherwise exposed to the risk associated with changes in volatility.

With regard to this situation, it is important to know precisely whether return variance constitutes an independent risk factor that is priced by the market or whether the returns on financial instruments and strategies that are targeted at capturing the variance risk premium can be explained by other means. For one thing, a profound knowledge about

the risks to which an asset offers exposure is necessary in the context of asset pricing to determine expected returns and asset prices that are commensurate to their risk exposure. Furthermore, this knowledge is of equal importance for related applications such as performance evaluations and attributions of institutional investors who engage in volatility related strategies and for risk management purposes since an assessment of the risks to which one is exposed arguably allows more efficient and successful hedging.

With reference to this, the objectives of this work are twofold. The first objective is to extract the variance risk premium from a set of traded options and to quantify its magnitude. Despite the fact that this is an approach taken by several other studies before (e.g. [Bakshi and Kapadia \(2003\)](#)), doing so can nevertheless provide interesting insights. Volatility related strategies are commonly regarded as extremely profitable, allowing to obtain Sharpe ratios that are substantially higher than simple equity investments (e.g. [Carr and Wu \(2009\)](#)). However, such strategies are often characterized by relatively frequent (small) gains but occasional losses of excessive magnitude ([Ilmanen \(2012\)](#)). Since unprecedented levels of volatility emanated during the period around and following the recent financial crisis and severely affected the returns on instruments whose value depends on return volatility, the dataset applied in this work offers interesting insights about how such strategies performed during this period and how average returns are affected over the long-run.

The second objective is to assess whether the extracted variance risk premium indeed represents compensation for the risk associated with a fluctuating return variance or whether the returns on such strategies can be explained by commonly applied asset pricing models.

1.2. Course of examination

For this work, I use an extensive set of options data on three stock indices from the United States (US) and five from Europe as well as on 29 single names from the US over the period from January 1996 until August 2015. In order to quantify the variance risk premium on each of these underlyings, I follow the procedure proposed by Carr and Wu (2009) and use the notion of a variance swap contract which is effectively a forward contract that pays the difference between the actually realized return variance over a predefined time period and a variance strike. The variance strike in this contract equals the expected risk-neutral variance and can be synthesized from a set of European options using a relatively model-free approach. Even though the estimate obtained with this method is subject to an approximation error when the asset price process is not purely continuous, it offers the advantage of a theoretical basis and does not impose further structural restrictions on the asset price process, which would inevitably be the case if one tried to calibrate a certain model to real data in order to derive the variance risk premium.

If the market does not require a non-zero risk premium that is embedded in the prices of traded options, the variance strike should equal the ex-post realized variance on average. As a consequence, it is possible to quantify the average variance risk premium as the time series average of the difference between the ex-post realized variance and the risk-neutral expected variance over the corresponding time period, i.e. the variance strike, or alternatively as the average return on a variance swap contract.

Since the results suggest that variance risk premia are indeed significantly different from zero and negative for all indices and the majority of single names, but differ substantially in their absolute magnitude, the analysis is continued with an attempt to link the variance risk premium, or variance swap returns, to the underlying's exposure to systematic variance. For this purpose, I compute a variance beta that measures the covariation between the return variance of a proxy for the market portfolio and that of the underlying under consideration. Variance swap returns are then regressed on this variance beta. Because this analysis reveals that the variance beta can only explain variance swap returns for the US indices but not for single stocks, it is further examined whether variance swap returns can be explained by a systematic variance risk factor that is proxied by the return on a variance swap with the S&P 500 as underlying. The underlying rationale for this course of action is that different studies find evidence that idiosyncratic return variances have a tendency to move together and possibly exhibit a common factor structure (Herskovic et al. (2014)) to which a diversified index, however, should not be exposed. Investors may require compensation for the risk associated with common

movements in idiosyncratic return variances and this compensation may contribute to the observed variance risk premium (Schürhoff and Ziegler (2011), Gourier (2015)).

In order to further investigate whether return variance constitutes an independently priced risk factor, it is examined whether the classical Capital Asset Pricing Model (CAPM) or the Fama and French (1993) three-factor model are able to explain variance swap returns. Even though both models generate significantly negative regression betas with respect to market excess returns that are consistent with the commonly observed negative correlation between equity returns and return variance (Glosten et al. (1993)), regression alphas are mostly negative and significantly different from zero, which is suggestive of one or more additional priced risk factors.

Because the total return variance reflects the continuous, i.e. pathwise variation as well as discontinuous price movements or jumps, I attempt to examine whether the risk associated with each of these types of price variation separately commands compensation that contributes to the observed variance risk premium. For this purpose I construct two risk-factor-mimicking portfolios from traded options that are targeted at offering exposure to one risk while being relatively unaffected by the other. Since the validity of any conclusion about the extent to which jump risk is priced crucially depends on the ability of the jump-risk-mimicking portfolio to reliably capture the exposure to jumps, it is tested whether positive returns on this factor coincide with jumps detected by the non-parametric jump detection test of Lee and Mykland (2008), which appears to be the case. Even though both constructed risk factors significantly contribute to explaining variance swap returns and continue to be significant when commonly used risk factors are included as control variables, abnormal returns remain mostly negative and significantly different from zero. Due to this, I consider the effect of the chosen specification of variance swap returns. In the initial setting, regressions are performed using continuously compounded variance swap returns to account for the substantial skewness and kurtosis in the distribution of raw returns. However, this has the drawback of shifting the mean return further into the negative domain, which may contribute to the persistently negative alphas. Due to this, robustness tests with raw variance swap returns instead of continuously compounded returns are performed and suggest that the specification has a certain impact on results but abnormal returns, especially for the US indices, often remain significant.

Finally, motivated by a relatively recent study by Drechsler (2013) who finds that model uncertainty can help substantially to explain the magnitude of the observed variance risk premium, I examine the relation between a proxy for model uncertainty and variance swap returns.

2. Literature review

For one thing, this work is related to prior studies that examine whether a variance risk premium exists and in particular to studies that examine the existence of the variance risk

premium not only for stock indices but also for single stocks. One of the earlier studies that provide empirical evidence for the existence of a negative variance risk premium in index options is the work by [Bakshi and Kapadia \(2003\)](#) in which a long call option is dynamically hedged with the underlying. Since this strategy is effectively market-neutral and should therefore only be exposed to volatility risk, [Bakshi and Kapadia \(2003\)](#) interpret the negative average returns on the strategy as evidence for the existence of a negative volatility risk premium. While the outlined approach allows to infer the variance risk premium rather indirectly, [Carr and Wu \(2009\)](#) are among the first who use a model-free approach to synthesize variance swap rates from a sample of traded options on five US stock indices and 35 individual stocks in order to examine the existence and dynamics of the variance risk premium. This model-free approach allows a direct quantification of the variance risk premium which is defined as the difference between the risk-neutral expected and subsequently realized variance. Their results suggest that average variance risk premia are negative and substantial for the five indices whereas premia for individual stocks exhibit considerable variation, are not always significant and can even be positive. They further examine whether observed variance risk premia are related to a systematic variance factor proxied by the return variance of the S&P 500. This analysis reveals that assets with higher exposure to the variance factor are associated with more negative variance risk premia, which leads them to conclude that investors dislike elevated levels of market volatility and are willing to accept negative average returns on variance swaps to insure against rising market volatility. Moreover, [Carr and Wu \(2009\)](#) find that frequently used asset pricing models such as the Capital Asset Pricing Model or the [Fama and French \(1993\)](#) model are unable to explain variance swap returns.

[Driessen et al. \(2009\)](#) use the same model-free approach to derive variance risk premia embedded in S&P 100 stock index options and options on all its individual constituents. Similar to [Carr and Wu \(2009\)](#), they find a significantly negative variance risk premium for the S&P 100 whereas the variance risk premiums on individual stocks are often zero or even positive. Since an index variance risk premium should reflect the variance premia of all its constituents, they conclude that a negative variance risk premium for the index together with non-existent or even positive variance risk premia on its constituents is only reconcilable with priced correlation risk. In order to empirically test this hypothesis, they implement a trading strategy that is targeted at capturing the correlation risk premium and earns significant abnormal returns. Based on their results, [Driessen et al. \(2009\)](#) also link the expensiveness of index options relative to individual stock options to the insurance against undesirable marketwide correlation increases that index options offer but individual options do not.

[Schürhoff and Ziegler \(2011\)](#) also synthesize variance swap rates to decompose total variance risk and examine the separate pricing of systematic and idiosyncratic variance risk for the S&P 100 and NASDAQ 100 index and all index con-

stituents. Consistent with previous findings, their results suggest that systematic variance carries a negative risk premium. Moreover, they also find that common idiosyncratic variance risk, i.e. the risk of comovements in the variances of idiosyncratic stock returns, carries a risk premium which is, on average, positive. Since the total variance risk premium is the sum of the two components, [Schürhoff and Ziegler \(2011\)](#) argue that the non-existent variance risk premium on individual stocks documented by [Driessen et al. \(2009\)](#) is due to the fact that the two components offset each other for S&P 100 constituents. They further attribute the positive risk premium on common idiosyncratic variance risk to financial intermediaries that are net-long in options on individual stocks. Since these intermediaries are not able to perfectly hedge options and are exposed to the associated idiosyncratic variance risk, they require compensation for upholding this position. Moreover, with regard to the results by [Driessen et al. \(2009\)](#), [Schürhoff and Ziegler \(2011\)](#) argue that the return on dispersion trading strategies that target to capture the correlation risk premium does not represent compensation for the pure risk of marketwide increases in correlation but rather reflects a combination of risk premia on systematic variance and common idiosyncratic variance risk.

In a related study, [Gourier \(2015\)](#) introduces an affine jump-diffusion model that considers the factor structure of asset returns as well as that of idiosyncratic return variance, in which variance swap rates and the variance risk premium – using the same definition as [Carr and Wu \(2009\)](#) – can be derived in closed-form. Further analyses show that model-based variance swap rates match synthetic variance swap rates remarkably well. [Gourier \(2015\)](#) finds a negative variance risk premium for all stocks that rises in absolute magnitude when the time to maturity increases. In contrast to [Schürhoff and Ziegler \(2011\)](#), she further finds that idiosyncratic variance risk carries a negative risk premium whose contribution to the overall variance risk premium is substantial and amounts to 80% on average.

The presented selection of studies shows that the existence of a negative variance risk premium for stock indices is well documented whereas the results and in particular the sign of the variance risk premium for single stocks are more ambiguous.

Apart from studies that document the existence of a variance risk premium, this work is also related to studies that attempt to examine whether the wedge between risk-neutral expected and actually realized variances predominantly reflects the risk that arises from stochastic volatility or rather results from compensation for the risk of discontinuous price movements, i.e. price jumps. For instance, [Todorov \(2010\)](#) examines the dynamics of the variance risk premium over time and especially how it is related to price jumps. For this purpose he fits a general semiparametric stochastic volatility model to S&P 500 data and uses deep out-of-the-money and close-to-maturity options to estimate the risk-neutral tail jump intensity. The chosen approach allows him to infer the contribution of stochastic volatility and jumps to the ex-ante variance risk premium which he measures as the difference

between a model-based estimate for the future expected variance under the physical measure and the VIX index. He concludes that jumps play a crucial role in explaining the variance risk premium. Since it increases after the occurrence of price jumps and reverts only slowly to its long-run mean thereafter whereas the impact of jumps on future market dynamics is limited, he concludes that investors' perception of jump risk is time-varying.

In a related work, [Bollerslev and Todorov \(2011\)](#) employ extreme value theory together with high frequency and options data in a nonparametric setting to decompose the equity risk premium and the variance risk premium into two separate components for diffusive and jump risk. Their results suggest that on average approximately 5% (in absolute terms) of the equity risk premium and more than half of the variance risk premium represent compensation for jump tail events.

3. Theoretical foundation of the variance risk premium

For a negative risk premium, such as the observed variance risk premium, to persistently occur, two complementary forces that are outlined in the following must coincide. First, the party requiring the risk premium must be exposed to essentially non-diversifiable, i.e. systematic, or non-hedgeable risk since otherwise a risk premium would not be justified and competitors would underbid each other until the market price of the asset equals its no-arbitrage price. The existence of such a non-diversifiable or non-hedgeable risk is strongly linked to the assumed price process of the underlying. It is well known that if the price of the underlying is assumed to follow a Geometric Brownian motion with constant return variance, as in the model by [Black and Scholes \(1973\)](#), markets are essentially complete and options are redundant securities because any option can be perfectly replicated by holding the commensurate amount of the underlying and the risk-free bond. The crucial aspect of such a setting is that the option's systematic risk is fully accounted for by the price of the underlying used in the replicating portfolio ([Hull \(2009\)](#)). Thus, in a world such as that described by [Black and Scholes \(1973\)](#), any exposure to options can be hedged perfectly and, if there is a competitive market for options, option prices should be solely determined by no-arbitrage conditions without the explicit consideration of risk premia. If stochastic volatility or random jumps in the asset price process are introduced, however, there are two additional variables that can change randomly and affect the price of an option. Consequently, it is no longer possible to perfectly replicate an option's payoff and thus hedge the option through simple trading in the underlying and the risk-free bond. As proven by [Bajeux-Besnainou and Rochet \(1996\)](#), introducing stochastic volatility that fluctuates independently of the price of the underlying (in contrast to a setting where volatility is a deterministic function of the asset price (e.g. [Dupire \(1994\)](#))) in a continuous time setting makes a classical European option always a non-redundant security. Because neither stochastic volatility nor jumps are traded assets,

the market is incomplete with respect to states in the world where these variables change ([Staum \(2007\)](#)). Technically, market incompleteness means that there is no unique equivalent martingale measure, or risk-neutral probability distribution, for which all discounted asset prices are martingales, but multiple such measures exist ([Bajeux-Besnainou and Rochet \(1996\)](#)). Due to the absence of assets that allow to trade these particular risks, it is also not possible to complement the underlying and the risk-free bond with an additional asset whose price exclusively depends on these variables and thereby replicate and hedge the option perfectly. As a consequence, the holder of an option position is left with the essentially non-hedgeable risk of random changes in the variance of the underlying or price jumps and associated changes in the value of her option. Thus, depending on whether these additional sources of risk are systematic and therefore correlated with aggregate consumption, they may command a risk premium. For instance, in the model proposed by [Hull and White \(1987\)](#), stochastic volatility does not carry a non-zero risk premium because it is explicitly assumed to be uncorrelated with aggregate consumption. Likewise, for example [Merton \(1976\)](#) assumes that the source of price jumps is company- or industry-specific information so that the contribution of price jumps to stock returns is non-systematic. However, there is considerable empirical evidence that return variances of different stocks tend to move together (e.g. [Andersen et al. \(2001\)](#)) and correlations in changes of implied volatilities make it difficult to eliminate vega risk by simple diversification ([Engle and Figlewski \(2014\)](#)). Thus, it can reasonably be assumed that variance risk is systematic and difficult to diversify. As a consequence, options will be priced to reflect the required compensation for random changes in the volatility (e.g. [Bates \(2000\)](#)). Similarly, [Ang and Chen \(2002\)](#) examine how correlations differ for upside and downside moves and find that correlation asymmetries are more pronounced for extreme downward moves, which can be interpreted as an indirect indication that jump risk may be systematic, thus potentially justifying a jump risk premium.

Note that for example in the model by [Heston \(1993\)](#), an additional option that is already traded would complete the market ([Staum \(2007\)](#)) and theoretically allow to perfectly hedge the risk of random changes in return variance. However, even in this situation, a perfect hedge still requires that the applied option pricing model be correct and the correct input parameters be used to reliably estimate each option's sensitivity with respect to random changes in the return variance. This leads to the closely related practical problem of model uncertainty. In this context [Broadie et al. \(2009\)](#) argue that market makers in option markets may require a risk premium since the necessary estimation of parameters such as the spot volatility, long-run mean levels of volatility and volatility mean reversion parameters and the associated determination of hedge ratios in the presence of stochastic volatility and jumps is subject to considerable estimation risk and may lead to the effect that option market makers cannot hedge their option exposure perfectly. For instance, [Green](#)

and Figlewski (1999) analyze the effect of inaccurate volatility estimates on delta-hedged short positions in index options by comparing the performance using as input to the option pricing model the best historical estimate of volatility and actually realized volatility over the remaining life of the option. Their results show that model risk due to an inaccurate volatility estimate has economically substantial effects even when positions are delta-hedged daily. Thus, the pricing of this estimation risk by risk-averse market makers could also contribute to the observed difference between option-implied and realized variances.

Empirical evidence that option market makers may indeed find it difficult to hedge their inventory is given by Gârleanu et al. (2009) who argue that option prices should not be affected by demand pressures when competitive intermediaries can hedge their option positions perfectly. They formalize in a model the notion that risk-averse intermediaries require compensation for their inability to hedge their exposure and link the amount of unhedgeable risk to the net-position these market makers hold. Their empirical results suggest that option expensiveness, defined as the difference between the average of options' implied volatility and a volatility forecast is indeed related to the net-position market makers have in options. In particular, index options in which intermediaries are net-short so that increases in volatility are associated with losses to the intermediary are visibly more expensive than individual equity options in which intermediaries are net-long. Further, Fournier and Jacobs (2015) also assume that market makers face unhedgeable risks and find that an option market makers' increasing inventory exposure to market variance risk, which is measured as the aggregate Black and Scholes (1973) vega, is associated with a significantly more negative variance risk premium (using the same definition of the variance risk premium as applied here). In this context, the results of Gârleanu et al. (2009) and Fournier and Jacobs (2015) offer an interesting explanation for the differential magnitude of variance risk premia in equity index options and options on individual stocks (e.g. Carr and Wu (2009), Driessen et al. (2009)) with regard to the net-position held by market makers in these contracts. Note, however, that the cited studies take the demand for options as exogenously given. Thus, they offer a potential explanation for why market makers may demand a variance risk premium but leave open the question why investors should be willing to pay it. This leads to the second prerequisite for the persistent existence of a variance risk premium.

While market makers seem to be exposed to unhedgeable systematic risk which justifies a risk premium, the second prerequisite for its persistent existence is that the assets on which an investor pays the premium, namely (index) options and variance swaps, offer insurance against an undesirable state of the world. If this was not the case, investors would not be willing to pay the substantial negative premium that is often found. A risk-averse investor is usually characterized by marginal utility that decreases with wealth. Due to the asymmetry of changes in utility that result from positive and negative changes in wealth or consumption of equal magnitude

with a concave utility function, a risk-averse individual will always be willing to give up wealth in the good state of the world, i.e. reduce current consumption, and pay insurance against realization of the bad state in which wealth is low and marginal utility is high. In this context, the apparent question is what exactly causes the risk-neutral expected variance to frequently exceed the realized variance. A possible explanation is a phenomenon called asymmetric volatility (e.g. Wu (2001)) that describes a negative correlation between market returns and return variance. Due to this negative correlation, variance swaps naturally offer insurance against substantial market declines, which may explain why investors are willing to pay a premium for such instruments, i.e. accept a variance strike that is too high relative to the variance one would expect to realize under the physical measure.

Moreover, a randomly changing return variance also implies that an investor's final wealth is not only risky but also ambiguous¹ or uncertain and there is empirical evidence of ambiguity aversion, i.e. that people prefer to act on known rather than on unknown probabilities (Ellsberg (1961)²).

An interesting and intuitive explanation how model uncertainty, which is closely linked to ambiguity aversion, can also help to explain risk-neutral expected variances that frequently exceed realized variances is offered by Drechsler (2013). He develops a model in which the representative investor has a reference model about the evolution of certain economic fundamentals but is not confident whether this model is indeed correct. To account for this situation and to derive decisions that are robust to model uncertainty, the investor considers alternative models that are statistically difficult to distinguish from the reference model and evaluates his decisions under the model which represents the "worst case". The investor is particularly concerned that the reference model underestimates the intensity and magnitude of potential jumps. Thus, the risk-neutral probabilities that are determined under this "worst case" model are tilted toward states of the world in which wealth is low and marginal utility is high, in particular to states in which large negative jumps in the expected growth rate of cash flows occur. As such important shocks to the economic state affect asset prices and thus return variance, this mechanism directly translates into a higher risk-neutral expected variance. Because, irrespective of the sign, the realized return variance is positively affected by the occurrence of jumps, a variance swap obviously offers insurance against such adverse movements. Moreover, an implication of this model is that the

¹Risk refers to a situation in which possible outcomes are uncertain but the distribution of outcomes is known. In contrast, ambiguity refers to a situation in which the outcome itself as well as the distribution of outcomes is not known (e.g. Anderson et al. (2009)).

²Ellsberg (1961) describes a situation in which people can choose a ball to be drawn from two different urns that contain red and black balls. A certain prize is received if a red ball is drawn and a smaller prize if a black ball is drawn. Urn I contains exactly 100 balls but nothing is known about the relative proportion of red and black balls. Urn II contains exactly fifty black and fifty red balls. Once asked to bet on the outcome that a red (black) ball will be drawn, the majority of people prefer Urn II (Urn I) over Urn I (Urn II).

magnitude of the variance risk premium depends on the degree of investors' uncertainty. A circumstance that gives support to this model is that, once calibrated, it is able to generate the variance risk premium with a comparatively low level of risk aversion and to capture additional features of asset returns.³ This is mainly a result of the explicit consideration of model uncertainty. In a related model by Drechsler and Yaron (2011) where this model uncertainty is not considered and the variance risk premium only arises from shocks to long-run consumption growth, larger jumps and a higher risk aversion are necessary to produce results that match important properties of the data. What both models (and many others) have in common, however, is that jumps are typically needed to produce variance risk premia that are close to the data estimates.

4. Methodology and data

4.1. Theoretical basis for extracting the risk-neutral expected variance

In this section, the methodology that is used to extract the risk-neutral expected variance from a set of traded options is outlined. The derivation of the risk-neutral expected integrated variance goes back to the work of Demeterfi et al. (1999), Britten-Jones and Neuberger (2000), Carr and Madan (1998) and others. Carr and Madan (1998) show that the integrated variance of the futures price process can be replicated through a static portfolio of European options across a continuum of strike prices and dynamic hedging in the underlying futures contract if continuous trading is possible, interest rates are constant and the underlying futures price is a continuous semi-martingale. Their solution builds on the work of Neuberger (1994) who finds that the payoff from delta-hedging a contract that pays the log of the spot price at maturity equals the difference between the realized variance and the variance assumed for hedging. Moreover, this log contract can be replicated through a static position in options (Breed and Litzenberger (1978)).

In particular Carr and Madan (1998) derive the following expression for the time-0 conditional expectation of variance of the futures price process over the time horizon $[0, T]$ under the risk-neutral probability measure \mathbb{Q} where F_t denotes the time t price of the futures contract that expires at time T and $RV_{t,T}$ denotes the annualized realized variance over the time horizon $[t, T]$.

$$E_0^{\mathbb{Q}}[RV_{0,T}] = E^{\mathbb{Q}}\left[\frac{2}{T}\left(\log\left(\frac{F_0}{F_T}\right) + \frac{F_T}{F_0} - 1\right) - \frac{2}{T}\int_0^T\left(\frac{1}{F_0} - \frac{1}{F_t}\right)dF_t\right] \quad (1)$$

³In particular, the model is able to generate the high equity risk premium, low risk-free rate, excess volatility of equity returns relative to fundamentals, a substantial variance risk premium that has predictive power for equity returns and a realistic implied volatility surface that captures the implied volatility skew for different maturities.

The second term on the right hand side of equation (1) represents the payoff of a continuously rebalanced position in the underlying futures contract. The first term represents a log contract that pays $f(F_T)$ at maturity. The final payoff of this log contract can be expressed as the final payoff of a static portfolio of European out-of-the-money (OTM) put and call options on the underlying futures contract⁴ across a continuum of strike prices K with expiration at time T where each option is inversely weighted by the square of its strike price K . Thus, the expression shown in equation (2) can be derived.

$$E_0^{\mathbb{Q}}[RV_{0,T}] = \frac{2}{T}E^{\mathbb{Q}}\left[\int_0^{F_0}\frac{1}{K^2}(K - F_T)^+dK + \int_{F_0}^{\infty}\frac{1}{K^2}(F_T - K)^+dK - \int_0^T\left(\frac{1}{F_0} - \frac{1}{F_t}\right)dF_t\right] \quad (2)$$

Note that the terminal option payoff in equation (2) is the payoff on European options on the futures contract. However, since the options and the futures contract expire at time T and the futures price has to equal the spot price at maturity to prevent arbitrage, European options on the futures and options on the spot are effectively equivalent. Moreover, in a risk-neutral world, the drift of the futures price is zero when the money market account is taken as numeraire implying that the futures price is a martingale (Carr and Madan (1998)). Thus, the expected value from continuous trading in the futures contract is zero so that the last term in the square brackets can be omitted.

Since the current price of any asset should equal the present value of future expected cash flows, the final payoff of the option portfolio can be rewritten as the terminal value of the current price of the option portfolio. Thus, the time t conditional expectation of variance over a time horizon $[t, T]$ under the risk-neutral measure can be stated as

$$E_t^{\mathbb{Q}}[RV_{t,T}] = e^{r_{t,T}(T-t)}\frac{2}{T-t}\left[\int_0^{F_0}\frac{1}{K^2}P_t(K, T)dK + \int_{F_0}^{\infty}\frac{1}{K^2}C_t(K, T)dK\right], \quad (3)$$

where $r_{t,T}$ is the time- t risk-free rate over the time horizon $[t, T]$ and $P_t(K; T)$ and $C_t(K, T)$ are the time- t prices of European put and call options on the spot with strike price K that expire at time T . Equation (3) also denotes the variance

⁴Carr and Madan (1998) formally show that any twice differentiable payoff function, $f(F_T)$, of the terminal futures price can be re-written as $f(F_T) = f(\kappa) + f'(\kappa)[(F_T - \kappa)^+ - (\kappa - F_T)^+] + \int_0^{\kappa}f''(K)(K - F_T)^+dK + \int_{\kappa}^{\infty}f''(K)(F_T - K)^+dK$. Setting the arbitrary parameter κ equal to F_0 allows to derive the expression shown in equation (2).

strike in a variance swap contract that is initiated at time t and expires at time T .

Equation (3) measures the risk-neutral expected variance exactly when the price of the underlying evolves according to a purely continuous price process without jumps. However, Jiang and Tian (2005) show that this expression also yields an accurate estimate under more realistic assumptions when the asset price process is allowed to exhibit jumps and the effect of implied jumps is included in the above measure so that an estimate of the total quadratic variation can be obtained. Further, Carr and Wu (2009) show in a simulation study that the jump-induced error is typically very small. In general, the approximation error due to jumps, i.e. the difference between the “true” risk-neutral expected variance and the estimate obtained from equation (3) will be positive when negative jumps dominate and can become significant when jumps significantly contribute to overall volatility, presumably in times of stress (Du and Kapadia (2011)).

Moreover note that equation (3) theoretically captures the risk-neutral integrated variance of the futures price process. However, under the assumption of deterministic interest rates and dividend yields, the spot and futures price should have the same quadratic variation. Moreover, since interest rates are at historically low levels for a substantial period covered by the dataset and the time horizon over which the risk-neutral expected variance is approximated is relatively short (22 trading days), a violation of these assumptions would probably have no significant effect on the results. Following the convention to determine the payoff on a variance swap that is described by Ait-Sahalia et al. (2015b), the annualized realized variance over the 22 trading day horizon $[t, t+22]$, where t now denotes a specific trading day, is defined as the annualized sum of daily squared log-returns shown in equation (4) based on a day-count-convention of 255 business days per year. The spot price used in equation (4) is adjusted for stock splits and spin-offs.

$$RV_{t,t+22} = \frac{255}{22} \sum_{i=1}^{22} \left(\log \left(\frac{S_{t+i}}{S_{t+i-1}} \right) \right)^2 \quad (4)$$

As noted by Bollerslev et al. (2011), the use of model-free realized volatilities, computed by summing squared returns from high-frequency data, generally allows a more accurate ex post observation of historical volatility than the expression used in equation (4) and would therefore naturally lend itself for an accurate measurement of the variance risk premium as defined below. However, in an earlier version of their study, Bollerslev et al. (2008) find, based on the model of Heston (1993), that the mean bias in the volatility risk premium using model-free implied volatility and realized volatility estimated from daily returns was only about 1.05% of the theoretical volatility premium for a sample size of 600.⁵ I consider this magnitude of bias sufficiently small to justify

the use of equation (4) to measure the realized variance over the relevant horizon.

4.2. The variance risk premium

Following Carr and Wu (2009), the variance risk premium is defined as the difference between the actually realized variance over the time horizon $[t, T]$ and the time- t expected variance under the risk-neutral measure over that same time period, i.e. $VRP_{t,T} = RV_{t,T} - E_t^Q[RV_{t,T}]$. In this sense, $VRP_{t,T} \cdot 100$ is the payoff in monetary units to an investor who holds a long position in a variance swap contract with notional 100 that is initiated at time t and expires at time T , whereas $RVRP_{t,T} = \frac{RV_{t,T}}{E_t^Q[RV_{t,T}]} - 1$ is the excess return to the investor when the variance swap rate is thought to be the initial investment. $LVRP_{t,T} = \log(RV_{t,T}/E_t^Q[RV_{t,T}])$ can therefore be regarded as the continuously compounded excess return on this variance swap contract.

When no confusion arises, I drop subscripts in the text. Specifically, I use VRP, RVRP and LVRP to refer to $VRP_{t,T}$, $RVRP_{t,T}$ and $LVRP_{t,T}$ in the text.

4.3. Data

The applied dataset includes daily option price information on three US and five European stock indices as well as 29 single stocks from the US. A detailed overview of the included indices and stocks along with further information is shown in Table 1 below.

Daily settlement as well as bid and ask prices for options on European indices as well as the corresponding index series and all overnight indexed swap (OIS) rates are retrieved from Thomson Reuters Datastream. Bid and ask prices of options on US indices and stocks, the corresponding prices of the underlyings and the London Interbank Offered Rate (LIBOR) curve are retrieved from OptionMetrics. For the European indices, OIS rates are used as proxy for the risk-free rate. For the US underlyings, LIBOR rates are used as proxies for the risk-free rate prior to 2007 whereas from the year 2007 on, OIS rates are used as a proxy. The reason for this change is that derivative dealers generally used LIBOR as a proxy for the risk-free rate prior to 2007 but have switched to OIS rates for collateralized transactions in later years as a consequence of substantial rises in LIBOR rates during the financial crisis (Hull and White (2015)). Since the data used here stem from options traded on organized exchanges with central clearing authorities and collateral requirements, the application of these rates appears to be a reasonable choice.

If available, option prices for subsequent analyses are defined as the average of bid and ask prices. Otherwise, the daily settlement price is used. To obtain the final data sample, several exclusionary criteria are applied to the initial set of options. First, if applicable bid prices are required to be strictly positive and bid-ask-spreads to be equal or greater than zero. Second, options with less than seven days to maturity are excluded since prices of such options might be biased by liquidity and microstructure concerns

⁵This can be seen by comparing the figures in Table 1 of their paper. The mean bias in the volatility premium is about 0.0021 compared to a theoretical premium of -0.20.

Table 1: Data description

Note: Entries list the name of included stocks and indices as well as their tickers which are sometimes used in other tables to reference the respective underlying. Start and End denote the first and last day in the sample on which a valid estimate of the variance risk premium could be obtained. NOptions denotes the number of option quotes available after filters are applied. NVSwap denotes the number of days on which a valid estimate for the VRP could be obtained. NKLow and NKHigh denotes the average daily number of options with expiration date T_L and T_H , respectively, from which implied volatilities could be derived and interpolated. TimeGap denotes the average distance in days between T_L and T_H over which the risk-neutral expected variance is interpolated. A description of T_L and T_H is given in section 4.4.

Underlying	Ticker	Start	End	NOptions	NVSwap	NKLow	NKHigh	Time Gap
AEX Index	AEX	06/06/2011	29/07/2015	54692	736	22.66	21.92	19.91
CAC40 Index	CAC	04/01/2010	29/07/2015	66668	967	23.95	18.84	19.93
DAX Index	DAX	18/04/2006	29/07/2015	257088	1660	44.43	45.23	19.73
Dow Jones Industrial Average	DJX	31/12/2002	30/07/2014	95091	2061	15.40	13.93	19.58
Euro Stoxx 50	ESX	03/01/2011	29/07/2015	114076	813	40.54	40.96	19.69
NASDAQ100	NDX	04/01/1996	30/07/2014	304070	3278	31.26	25.01	18.60
SMI	SMI	06/06/2011	29/07/2015	48250	717	22.97	18.99	19.49
S&P500 Index	SPX	04/01/1996	30/07/2014	649410	3598	47.88	42.58	16.43
Alcoa	AA	19/08/1997	30/07/2014	35747	803	5.13	6.56	18.46
Altria (Philip Morris)	MO	05/01/1996	30/07/2014	55143	2163	5.29	6.42	18.35
Amazon	AMZN	15/12/1997	30/07/2014	142691	2315	13.80	16.42	17.28
American Express	AXP	26/01/1996	30/07/2014	65397	1949	7.77	8.82	17.93
Amgen	AMGN	17/01/1996	30/07/2014	68730	2328	6.93	7.97	17.98
Analog Devices	ADI	20/08/1996	20/05/2014	29604	685	5.74	5.95	18.05
Apple	AAPL	04/01/1996	30/07/2014	272051	2392	23.22	28.25	16.81
Bank of America	BAC	17/01/1996	30/07/2014	51899	1467	5.22	6.40	15.60
Boeing	BA	27/02/1996	30/07/2014	77177	2032	7.77	8.83	17.24
Cisco	CSCO	04/01/1996	30/07/2014	63305	1932	6.23	7.84	17.76
Exxon Mobil	XOM	19/02/1997	30/07/2014	63227	1738	7.01	8.50	17.35
Facebook	FB	29/05/2012	30/07/2014	57679	501	20.35	21.38	8.22
General Electric	GE	17/01/1996	30/07/2014	60886	1860	5.63	7.18	17.14
Home Depot	HD	17/01/1996	30/07/2014	62245	1789	6.48	7.84	17.55
IBM	IBM	04/01/1996	30/07/2014	101647	2920	8.59	10.30	17.40
Johnson & Johnson	JNJ	26/01/1996	30/07/2014	47892	1238	5.92	6.88	16.92
McDonald's	MCD	29/02/1996	30/07/2014	56254	1406	7.01	8.13	16.49
Merck	MRK	22/02/1996	30/07/2014	62784	2038	6.38	7.60	17.96
Metlife	MET	13/02/2007	30/07/2014	37523	945	9.65	10.79	19.32
Microsoft	MSFT	04/01/1996	30/07/2014	87564	2441	7.78	9.91	18.13
Monsanto	MON	19/03/1996	30/07/2014	43459	1295	8.59	9.81	19.27
Nike	NKE	01/02/1996	30/07/2014	50464	1448	6.41	7.02	16.57
Pfizer	PFE	23/10/1996	30/07/2014	54150	1442	5.46	6.62	16.62
Procter & Gamble	PG	17/07/1996	30/07/2014	53938	1609	6.06	6.99	16.94
Starbucks	SBUX	18/09/1996	30/07/2014	53567	1530	7.06	8.41	18.92
Tesla	TSLA	21/07/2010	30/07/2014	64480	732	18.49	18.47	14.58
Valero	VLO	04/06/2001	30/07/2014	64166	1747	8.31	10.37	18.02
Verizon	VZ	25/03/1998	30/07/2014	59440	1399	7.12	8.69	16.65
WalMart	WMT	22/10/1997	30/07/2014	58038	1882	5.92	7.09	17.39

(Jiang and Tian (2005)). Third, index options with zero trading volume are excluded from the dataset because the prices of such options may not reflect true value (Jiang and Tian (2005)). Since most single stock options are less frequently traded than index options, a less strict criterion is applied to such options in that only options with zero trading volume and open interest smaller than 100 contracts are excluded.

Since the following applications require a dividend yield for the indices of which a reliable estimate is difficult to obtain, the option-implied dividend yield is obtained through a combined application of the put-call parity and spot-futures parity. In a first step, the time- t option-implied price, F_t^i , of

the futures contract that expires at time T is inferred through put-call-parity (e.g. Hull (2009)) that is shown in equation (5).

$$F_t^i = K + e^{r_{t,T}(T-t)} \cdot (C_t(K, T) - P_t(K, T)) \quad (5)$$

In order to select the put and call option used in equation (5), the methodology applied by the Chicago Board Options Exchange for calculating its volatility index (CBOE (2015)) is used. For every maturity on a given day, the pair of options with the same strike price and maturity for which the absolute price difference is smallest is used in equation (5). The option-implied forward looking dividend yield at time t

over the time horizon $[t, T]$, $\theta_{t,T}$, is then extracted through the spot-futures parity (e.g. Hull (2009)) that is shown in equation (6).

$$F_t^i = S_t e^{(r_{t,T} - \theta_{t,T})(T-t)} \quad (6)$$

The outlined procedure ensures that the option pair used in the estimation of the implied dividend yield consists of at-the-money (ATM) options which are, especially for the indices considered here, typically relatively liquid. Thus, prices of such options and derived implied dividend yields can generally be expected to be reliable. Nevertheless, option-implied dividend yields are occasionally negative. Such negative values occur remarkably often during times of financial turmoil, especially during October 2008 and November 2011. Due to this, the negative yields most likely capture the effect of a discount rate used by market participants that exceeds the rate that is used to derive the implied dividend yield. Even though these negative dividend yields ensure that put-call parity is technically satisfied, it is obviously implausible to use them for further applications. Thus, the following adjustments are made to replace negative values: whenever possible, a replacement value for a negative forward looking dividend yield at a given maturity is obtained through linear interpolation or extrapolation of implied dividend yields at adjacent maturities on the same day. If a linear interpolation or extrapolation is not possible because the number of positive implied dividend yields at different maturities is smaller than two on a given day, the implied dividend yield for the same expiration date on the previous day is used as a replacement value. If none of these two adjustments results in a positive value or no value can be found, the implied dividend yield is set to zero. For those observations where the implied dividend yield is replaced, a new forward price corresponding to the adjusted implied dividend yield is computed.

For stocks a similar procedure is applied with the exception that the implied present value of dividends is derived instead of an implied dividend yield.

After the derivation of implied dividend yields and present values of dividends, a fourth exclusionary criterion is applied before implied volatilities are derived. As outlined by Ait-Sahalia and Lo (1998), in-the-money (ITM) options are traded relatively infrequently compared to ATM or OTM options. Therefore, prices of such options, and implied volatilities derived from these prices, tend to be unreliable. For this reason, ITM options are removed from the dataset.

For European options, implied volatilities are inferred based on the model of Black and Scholes (1973) (henceforth referred to as B-S model).⁶ For American options, implied volatilities provided by OptionMetrics are used. For such options, OptionMetrics uses a binomial tree approach accounting for the effect of the early exercise premium.

⁶This is done with an adjusted version of Mark Whirly's code "Fast Matrixwise Black-Scholes Implied Volatility" for Matlab that is available under the following address: <http://www.mathworks.com/matlabcentral/fileexchange/41473-fast-matrixwise-black-scholes-impli-ed-volatility>.

Index options with implied volatilities greater than 80% and single stock options with implied volatilities greater than 100% are removed from the dataset. Such observations are probably outliers that would corrupt the implied volatility surface.

In addition to the options data, price information from the Center for Research in Security Prices, factor returns from Kenneth French's website and data from the Survey of Professional Forecasters, which are introduced in detail when applied, are used.

4.4. Implementation of the method to extract the risk-neutral expected variance

In order to obtain an estimate of the risk-neutral expected variance over the time horizon $[t, T]$, equation (3) is numerically evaluated using the trapezoidal method. A practical problem arises since equation (3) requires an infinite number of options across a continuum of strike prices whereas the number of traded options is finite. To cope with this situation, I generate 5000 artificial options that expire at the two expiration dates T_L and T_U closest below and above T over an equally-spaced range of strike prices of ± 8 standard deviations from the current spot price. This standard deviation is estimated as the average of the implied volatilities of the two options that are closest to being at-the-money and expire at time T_L . In a simulation based on a model with stochastic volatility and jumps, Jiang and Tian (2005) show that truncation errors are generally negligible if the truncation points, i.e. the highest and lowest observed strike prices, are more than two standard deviations from the current forward price, F_0 , and can be further reduced by applying the extrapolation scheme outlined below. They further show that discretization errors resulting from non-continuous strike prices are negligible when the gap between consecutive strike prices is smaller than or equal to 0.35 standard deviations. Thus, 5000 options over a range of strike prices between ± 8 standard deviations from the current spot should help to alleviate truncation and discretization errors.

In order to generate the artificial option prices, implied volatilities are needed. As noted in section 4.3, implied volatilities for European options are obtained based on the model by Black and Scholes (1973) whereas implied volatilities provided by OptionMetrics are used for American options. For any given day t and the two expiration dates T_L and T_U , implied volatilities are then interpolated across moneyness, defined as $k \equiv \log(K/F_t^i)$, between the highest and lowest observed strike price using cubic splines. Jiang and Tian (2007) argue that the use of cubic splines leads to the convenient property that the implied volatility function is smooth over the range of observed strike prices, which is a direct implication of no-arbitrage constraints (e.g. Breeden and Litzenberger (1978)). For strike prices below the smallest and above the highest observed strike price, implied volatility is held constant at the value observed at these end points.⁷

⁷Jiang and Tian (2007) use a slightly different but similar linear extrapo-

The prices of artificial European options that expire at time T_L and T_U over the continuum of strike prices are obtained by inserting the interpolated implied volatilities into the B-S-formula. In the next step these option prices are used to numerically evaluate equation (3) and obtain an estimate of the risk-neutral expected variances over the two time horizons $[t, T_L]$ and $[t, T_U]$. An estimate of the risk-neutral expected variance over the desired horizon $[t, T]$, can then be obtained through linear interpolation between the expected risk-neutral variances over the time horizons $[t, T_L]$ and $[t, T_U]$. Note that, in order to ensure that the interpolation and extrapolation scheme can work properly and derived estimates of the risk-neutral expected variance are reliable, this procedure is only applied on days on which at least four traded options are available. Moreover note that the B-S model is only used as a convenient way to derive implied volatilities from and translate implied volatilities into options prices. In particular, it is not assumed that this model is a true representation of reality.

A description of the average time over which the risk-neutral expected variance is interpolated along with the average number of options that expire at times T_L and T_H is shown in Table 1.

5. Empirical analysis

5.1. Realized variance risk premia

Figure 1 plots the daily time series of annualized risk-neutral expected and actually realized volatilities over the subsequent 22 trading days, the variance risk premium, defined as $VRP_{t,T}$, and the index level for the S&P 500 index between January 4th 1996 and July 30th 2014. From Figure 1 it is evident that the approximation of the expected variance captures the dynamic of the subsequently realized variance quite well and that substantial changes in the variance risk premium go hand in hand with significant fluctuations in the index level. This is also confirmed by the high correlation coefficient of 0.6437 between risk-neutral expected and subsequently realized variance for the S&P 500 index. At the same time, however, the wedge between the two variance measures is also apparent. For most of the time, the variance risk premium is negative and relatively stable but experiences substantial fluctuations when major market changes occur and reaches unprecedented levels during the financial crisis.

Table 2 shows summary statistics for the variance risk premium over a 22 trading day horizon in three different forms, defined either as the time series average of $VRP_{t,T}$

lination scheme where the slope of the extrapolated segment is set equal to the slope of the interior segment at the endpoints. Even though this procedure better captures the observed skew in implied volatility surfaces for certain strike sections, it can lead to the implausible drawback that further-out-of-the-money call options have higher artificial prices than nearer-out-of-the-money call options, which occasionally happens here. Thus, the constant extrapolation scheme with zero slope is applied.

100, $LVRP_{t,T}$ or $RVRP_{t,T}$. The reason for introducing a logarithmic version of the variance risk premium is that the VRP and RVRP distributions exhibit substantial skewness and kurtosis, indicating that the two stem from a highly non-Gaussian distribution, while LVRPs generally appear more normal which might be beneficial for subsequent regression analyses.

Consistent with previous studies, the mean VRP is negative on all stock indices and statistically significantly different from zero at the 1%-level for the S&P 500, the AEX, and the Euro Stoxx 50 and significant at the 5%-level for the Dow Jones and NASDAQ 100. Only for the Swiss SMI and the DAX, the variance risk premium is not significant on any conservative level. The distributions of both, VRP and RVRP show substantial kurtosis and positive skewness and are reflective of high and positive returns which fatten the right tail of the distributions. In contrast, LVRPs exhibit a considerably lower skewness and kurtosis, which is due to the fact that especially the excessively high, positive returns during the recent financial crisis are alleviated through the logarithmic transformation. In this context, it is also noticeable that the standard deviations of the VRPs are excessive, especially for single stocks, which partly explains why the VRP is statistically different from zero only for comparatively few stocks. Indeed, even though VRP is negative in most cases, the null hypothesis that realized and expected risk-neutral variances do not differ on average cannot be rejected at the 5% significance level for 21 of the 29 stocks.

Mean LVRPs are negative for all underlyings and t-statistics⁸ are generally higher. The mean LVRP is significantly different from zero at the 1%-level for all underlyings. However, as pointed out by Driessen et al. (2009), a necessary remark in this context is that due to the concavity of the logarithmic function, Jensen's inequality implies that mean log variance risk premia are negative even under the null hypothesis of equality between risk-neutral expected and subsequently realized variance. When raw variance swap returns ($RVRP_{t,T} = \frac{RV_{t,T}}{E[RV_{t,T}]} - 1$) are used instead, mean returns are generally smaller in absolute magnitude than for the LVRPs but still negative for 35 underlyings. However the mean raw return is statistically significantly different from zero at the 5%-level for only 14 underlyings.

Qualitatively, the results in the left panel of Table 2 are similar to those obtained by Carr and Wu (2009) who find significantly negative variance risk premia for stock indices but only comparatively few single stocks for which the premia are significant. Despite this tendency, the results are quite different from those of Driessen et al. (2009) who find a significantly negative variance risk premium for the S&P 100 index but no evidence for the presence of a generally negative variance risk premium in individual constituents. In-

⁸Except for equation (9), all t-statistics and regressions, as well as standard deviations adjusted according to the method of Newey and West (1987) are obtained using Kevin Sheppard's MFE Toolbox for Matlab throughout this work. The toolbox is available under http://www.kevinsheppard.com/MFE_Toolbox.




Figure 1: Time series of risk-neutral expected and realized volatility, variance risk premium and the S&P 500

Note: The first plot shows the time series of annualized risk-neutral expected and subsequently realized volatility over 22 trading days in the period from January 4th 1996 to July 30th 2014 for the S&P 500 Index. Both measures are annualized on a basis of 255 trading days in a year. The second plot shows the realized variance risk premium over 22 trading days, $VRP_{0,T}$, over the corresponding sample period. The red horizontal line indicates a VRP equal to zero. The third plot shows the index level of the S&P 500 index over the same sample period.

Table 2: Summary statistics of variance risk premia

Note: Entries refer to the time period from January 4th 1996 until July 29th 2015. RV and EV denote the realized and risk-neutral expected variance, respectively. Realized volatility is calculated from daily returns and annualized by $\sqrt{255/22}$. t-stat denotes the relevant t-statistics adjusted for serial correlation according to the method by Newey and West (1987) with a lag length of 22 days. Sample std., Skew, and Kurt denotes the sample standard deviation, skewness and kurtosis. Sharpe ratio is the annualized Sharpe ratio calculated as the annualized mean IVRP or RVRP divided by the corresponding sample standard deviation adjusted for serial dependence according to Newey and West (1987) with 22 lags and annualized by $\sqrt{255/22}$.

Ticker	Realized volatility	VRRP as $(RV - EV) \cdot 100$					IVRP as $\ln(RV/EV)$					RVRP as $RV/EV-1$							
		Mean	t-stat	Sample std.	Skew	Kurt	Mean	t-stat	Sample std.	Skew	Kurt	Sharpe Ratio	Mean	t-stat	Sample std.	Skew	Kurt	Sharpe Ratio	N
AEX	16.12%	-0.68	-3.01	2.18	1.22	8.51	-0.32	-4.61	0.59	0.43	3.27	0.58	-0.12	-1.70	0.65	2.56	11.20	0.11	736
CAC	19.88%	-0.62	-1.93	3.38	2.30	12.56	-0.27	-4.61	0.57	0.45	3.40	0.50	-0.09	-1.35	0.66	2.99	16.70	0.07	967
DAX	20.21%	-0.84	-1.65	5.78	4.02	37.84	-0.34	-6.46	0.61	0.50	3.85	0.54	-0.12	-1.91	0.74	3.51	19.47	0.06	1660
DJX	15.17%	-0.84	-1.99	4.95	4.92	51.04	-0.50	-10.72	0.58	0.74	5.04	0.80	-0.25	-4.63	0.67	4.46	28.23	0.14	2061
ESX	19.60%	-0.95	-3.02	2.88	1.30	7.73	-0.33	-5.05	0.55	0.27	2.96	0.60	-0.15	-2.60	0.53	1.91	7.70	0.18	813
NDX	26.05%	-1.17	-2.52	7.38	3.17	24.10	-0.33	-10.68	0.51	0.59	4.65	0.64	-0.17	-4.55	0.60	4.50	33.35	0.13	3278
SMI	14.05%	-0.24	-0.65	3.06	3.18	16.51	-0.33	-4.26	0.64	0.91	4.36	0.54	-0.08	-0.68	0.93	3.61	17.74	0.03	717
SPX	17.10%	-1.25	-3.65	5.15	5.72	60.03	-0.54	-15.05	0.57	0.79	5.07	0.85	-0.29	-6.61	0.67	5.61	46.08	0.14	3598
AA	40.18%	-0.37	-0.17	21.06	6.71	58.86	-0.22	-3.99	0.55	0.67	4.81	0.48	-0.04	-0.51	0.77	3.93	22.72	0.03	803
MO	25.68%	-1.16	-1.79	9.53	2.97	27.06	-0.37	-7.32	0.71	0.42	4.46	0.54	-0.07	-0.78	1.25	8.63	96.92	0.01	2163
AMZN	47.01%	1.65	1.01	24.60	2.74	14.05	-0.23	-5.31	0.62	0.45	3.91	0.38	-0.01	-0.23	0.86	4.11	27.72	0.01	2315
AMGN	32.06%	-1.54	-2.12	10.51	1.42	12.37	-0.30	-4.34	0.57	0.64	4.44	0.33	0.00	0.01	0.80	3.53	18.79	0.00	1949
ADI	51.22%	0.52	0.17	23.07	2.16	10.88	-0.19	-7.05	0.49	0.10	2.95	0.50	-0.12	-2.83	0.59	2.17	10.59	0.10	2328
AAPL	39.31%	-1.31	-0.63	39.51	14.80	242.03	-0.31	-8.64	0.54	0.80	3.35	0.41	-0.06	-0.94	0.53	2.11	9.26	0.07	685
BAC	39.00%	7.30	1.52	46.18	5.19	31.82	-0.19	-3.42	0.61	0.77	4.64	0.30	0.03	0.36	1.15	14.55	275.48	0.04	2392
BA	28.94%	-1.35	-2.45	7.92	2.96	27.29	-0.26	-7.18	0.51	0.26	4.00	0.54	-0.12	-3.16	0.56	3.80	29.38	0.14	2032
CSCO	39.36%	-0.11	-0.10	14.34	1.92	10.54	-0.24	-5.38	0.62	0.29	3.22	0.42	-0.04	-0.73	0.72	2.55	11.85	0.02	1932
XOM	22.53%	-0.33	-0.33	12.02	9.05	98.71	-0.31	-6.73	0.57	0.72	5.26	0.55	-0.11	-1.45	0.90	7.55	84.13	0.04	1738
FB	43.97%	-2.21	-0.68	19.40	1.88	8.67	-0.25	-2.60	0.60	0.88	4.22	0.40	-0.03	-0.23	0.85	3.19	13.74	0.01	501
GE	30.84%	-0.82	-1.03	10.46	2.38	17.44	-0.23	-5.83	0.54	0.59	4.45	0.46	-0.06	-1.12	0.73	4.26	29.99	0.04	1860
HD	30.48%	-0.80	-1.01	12.18	7.80	93.72	-0.32	-7.60	0.55	0.49	4.77	0.61	-0.13	-2.54	0.78	7.69	87.12	0.09	1789
IBM	26.61%	-1.32	-3.14	6.87	1.22	12.72	-0.34	-9.08	0.58	0.22	3.21	0.57	-0.15	-4.08	0.58	2.49	12.76	0.13	2920
JNJ	19.07%	-1.26	-2.61	5.81	2.82	23.23	-0.43	-7.77	0.64	0.21	3.97	0.75	-0.19	-3.35	0.72	5.24	44.45	0.16	1238
MCD	18.85%	-1.29	-3.28	4.68	1.67	26.08	-0.41	-9.26	0.50	0.69	3.82	0.84	-0.24	-5.43	0.50	3.07	17.35	0.30	1406
MRK	26.06%	-0.81	-1.47	8.06	2.71	22.74	-0.29	-6.57	0.63	0.32	2.98	0.50	-0.08	-1.67	0.70	2.37	10.97	0.06	2038
MET	34.99%	1.02	0.35	30.18	10.83	140.21	-0.22	-4.11	0.53	1.35	9.33	0.46	-0.02	-0.16	1.28	9.28	104.98	0.00	945
MSFT	32.20%	-0.81	-1.20	9.94	3.09	20.24	-0.24	-6.53	0.55	0.34	3.39	0.45	-0.08	-1.86	0.62	2.71	14.56	0.06	2441
MON	32.56%	-1.04	-0.81	13.47	5.17	42.87	-0.29	-6.24	0.50	0.53	4.07	0.59	-0.14	-2.68	0.55	2.83	13.40	0.14	1295
NKE	29.08%	-0.99	-1.25	9.93	2.60	16.39	-0.31	-6.18	0.61	0.41	3.36	0.59	-0.10	-1.77	0.70	2.86	13.87	0.07	1448
PFE	27.28%	-1.32	-3.05	6.10	1.90	13.01	-0.30	-6.50	0.56	-0.01	3.78	0.58	-0.13	-2.96	0.56	2.56	13.24	0.16	1442
PG	21.38%	-0.87	0.50	18.83	8.20	75.10	-0.39	-6.31	0.68	1.03	5.98	0.54	-0.06	-0.44	1.44	7.22	62.59	0.01	1609
SBUX	33.55%	-0.39	-0.43	12.89	4.39	37.57	-0.26	-5.59	0.56	0.51	3.28	0.49	-0.08	-1.67	0.65	2.77	14.57	0.07	1530
TSLA	54.35%	-5.68	-2.16	22.93	1.66	6.89	-0.32	-4.59	0.60	0.31	2.74	0.58	-0.12	-1.71	0.61	1.85	6.57	0.11	732
VLO	38.84%	0.12	0.08	17.15	5.42	44.52	-0.18	-4.17	0.52	0.48	3.68	0.34	-0.03	-0.62	0.62	2.68	13.15	0.02	1747
VZ	22.19%	-1.31	-2.11	6.85	3.76	40.91	-0.33	-8.23	0.49	0.41	3.72	0.75	-0.18	-4.31	0.50	3.38	22.50	0.25	1399
WMT	24.47%	-1.11	-2.31	6.02	1.09	12.29	-0.35	-8.19	0.54	0.32	3.30	0.64	-0.17	-4.20	0.54	2.56	12.94	0.18	1882

deed, the mean VRP is significantly negative at the 5%-level for only 7 of 129 (5%) stocks in their sample whereas this is the case for 9 of 29 stocks (31%) here. In this context, it is necessary to point out that the 29 stocks in this dataset have specifically been selected according to the number of price quotes provided by OptionMetrics and the number of valid price quotes following the criteria outlined above to allow the computation of a high number of estimates of the risk-neutral expected variance. Options that are comparatively illiquid and thus carry a substantial illiquidity premium should trade at lower prices, which would directly translate into lower implied volatilities. Applying the same procedure to derive the risk-neutral expected variance to such options should then naturally lead to lower expected risk-neutral variances and, holding realized variance constant, to less negative or even positive variance risk premia. Thus, the results here do not necessarily contradict the findings by [Driessen et al. \(2009\)](#) but could simply reflect different sample selection criteria and resulting differences in the sample composition.

When turning to the profitability of variance swap investments, the results show that the average continuously compounded return on shorting a 22-day variance swap on one of the major indices over the entire sample period is in the range of 27% in case of the CAC40 to 54% in case of the S&P 500. In this context, it is noticeable that the average returns on single stock variance swaps are lower than those on some indices, especially the US indices, but not substantially lower.

However, even though the variance premia are substantial, they are relatively low compared to the figures found by other studies. This circumstance is most obvious when one turns to the left panel of Table 2 where the payoff on a long variance swap for a notional of 100 currency units is shown. For instance, [Carr and Wu \(2009\)](#) find average payoffs of \$2.5 and above for the S&P 500, the Dow Jones and NASDAQ. The lower payoffs over the entire sample might in part be due to the fact that especially during the financial crisis in 2008 and 2009 and its aftermath, realized variance has frequently and substantially exceeded expected variance.⁹ This observation is also in line with the argument of [Carr and Lee \(2009\)](#) who point out that short positions in variance swaps led to significant and unprecedented losses especially during the final quarter of 2008. They further outline that market makers' difficulty to hedge their exposure to variance swaps, in particular for single names, led to a complete collapse of the single name variance swap market in 2009. For the sake of completeness, however, it is necessary to point out that the variance risk premia obtained here are generally smaller, in absolute terms, even over the same time period covered by [Carr and Wu \(2009\)](#).

⁹In addition, the smaller variance risk premiums found here compared to those found by [Carr and Wu \(2009\)](#) on the same data are due to the different interpolation schemes. While [Carr and Wu \(2009\)](#) use a linear interpolation scheme, the cubic spline interpolation applied here usually leads to lower implied volatilities and therefore lower option prices since the typical shape of implied volatilities across moneyness is convex. However, as shown below, the influence of the interpolation scheme is comparatively small.

In order to evaluate the profitability of variance swap investments with regard to their risk-return profile Table 2 also shows the annualized Sharpe Ratios. Since a variance swap contract is essentially a forward contract that does not tie up capital until maturity, the given variance swap returns directly represent excess returns and the Sharpe Ratios are simply computed as the annualized average return on the variance swap divided by the annualized return standard deviation which is again adjusted for serial correlation according to the method of [Newey and West \(1987\)](#). Considering the risk-return trade-off of shorting a 22-day variance swap, the high Sharpe Ratios for LVRPs of up to 0.85 for the S&P 500 index are relatively close to the estimates of [Carr and Wu \(2009\)](#) and suggest that variance swaps can be attractive investments. The premia that investors are obviously willing to pay for holding long variance swaps appear substantial compared to the risk to which the short side is exposed. However, when one considers the right panel in Table 2 for RVRPs, it is apparent that the high Sharpe Ratios for the LVRPs are mainly the result of the logarithmic transformation which not only alleviates the effect of large positive returns and thereby reduces the mean but also lowers the standard deviation. Both effects cause the Sharpe Ratios from shorting variance swaps to be substantially higher for LVRPs.

Of course, the extent to which the reported results are reflective of the profitability of actual variance swap investments crucially depends on how well the synthetic variance swap rates approximate the prices of actually traded instruments. However, at least for the S&P 500 index, [Ait-Sahalia et al. \(2015b\)](#) find that the methodology applied here leads to synthetic variance swaps rates for maturities of up to six months that are roughly in line with those of actually traded instruments.

As outlined before, the unprecedented levels of realized variance that manifested during the financial crisis also led to previously unseen returns on variance swaps. In order to evaluate the extent to which these returns affect the results reported in Table 2, Table A1 shows the same descriptive statistics for the time period from January 4th 1996 until December 31st 2007. For most underlyings the payoffs on a variance swap per 100 currency units notional is slightly higher in absolute terms than for the entire sample, whereas the continuously compounded returns as well as the raw returns are comparable and sometimes even lower in absolute terms than for the entire sample. The comparable mean returns over the two different samples can be explained by the fact that a long position in a variance swap earned substantial positive returns during the turmoil of 2008. However, after return variances had exploded, the variance swap rates increased accordingly and remained at their elevated levels even when realized variances reverted to lower levels so that average returns appear to be relatively unaffected over the entire sample period. This pattern is also apparent in the first plot in Figure 1. With regard to the relatively stable average variance swap returns, it is obvious that the Sharpe Ratios which are higher compared to those over the entire sample period for LVRPs as well as for RVRPs are predomi-

nantly driven by substantially lower standard deviations. Because the average returns are close to each other for the two samples, the sheer magnitude of variance swap returns that were realized during the financial crisis is disguised and it is not immediately apparent what could have led to the subsequent collapse of the single stock variance swap market mentioned by Carr and Lee (2009). In this context, it is helpful to consider the magnitude of variance swap returns only during this time period. The average maximum raw return on a long position in a 22-day single stock variance swap that was initiated during the final quarter of 2008 across all underlyings except the indices was a substantial 496.99%. For the S&P 500, the Dow Jones, and the NASDAQ, average 22-day returns on a long variance swap initiated during the last quarter of 2008 were 162%, 96%, and 107%. However, note that these figures are based on synthetic variance swap rates and especially when jumps significantly contribute to volatility, which was certainly the case at this time, the approximation error of the applied procedure to derive the risk-neutral expected variance can be high and the “true” risk-neutral quadratic variation is likely to be significantly underestimated (Du and Kapadia (2011)). Accordingly, actual variance swap rates would probably have been higher and realized variance swap returns lower. As a consequence, the stated figures can only serve as a rough indication of actual returns. Nevertheless, when considering the fact that market makers faced substantial problems in hedging their variance swap exposures at these times (Carr and Lee (2009)), the sheer magnitude of potential losses helps to understand why the single stock variance swap market could have broken down. All in all, it can be said that shorting variance swaps would have been a profitable strategy over the entire sample period despite the substantial losses during the financial crisis and its aftermath. However, when one considers raw returns, it is apparent that these losses have a deteriorating effect on Sharpe Ratios and these investments do not appear to be much more attractive than simple equity investments.

In order to ensure that the previously reported and following results are not substantially affected by the applied cubic spline interpolation scheme that is used to derive the risk-neutral expected variance and that has been chosen mostly independent from a guiding theory, a robustness test is performed. In order to do so, I recalculate the previously shown descriptive statistics based on a linear interpolation of implied volatilities in the range of traded strikes and apply the same extrapolation procedure as before. Results are shown in Table A2 in the Online-Appendix. As can be seen, the results under the two different interpolation procedures are almost identical. In general, the payoffs per 100 currency units of notional and the realized variance swap returns have a tendency to be slightly more negative when the linear interpolation scheme is used. This is due to the fact that the typical shape of implied volatilities across the applied moneyness measure is convex so that the cubic spline interpolation leads to lower interpolated implied volatilities on average. Therefore option prices are lower than under the linear interpolation scheme which then translates into lower estimates of

the risk-neutral expected variance. However, the deviations are small and can reasonably be expected not to significantly alter the results.

5.2. Explaining variance risk premia

5.2.1. Exposure to overall market variance

From the previous results, it is obvious that the gap between risk-neutral expected and subsequently realized variance is persistent and statistically significant for indices as well as individual stocks. Investors are apparently willing to pay a significant premium for holding long positions in variance swaps. Since standard asset pricing theory postulates that only systematic risk factors can command such premia (e.g. Ross (1976)), it is instructive to investigate which systematic factors help to explain variance swap returns.

Payoffs on variance swaps naturally depend on the return variance of the underlying asset. However, only that portion of the asset’s return variance that is systematic and cannot be diversified away in a portfolio of variance swaps can theoretically command a risk premium. Consequently, it is a reasonable starting point to test whether the covariation in an asset’s return variance with that of the market portfolio helps to explain variance swap returns. This covariation is measured by a variance beta for each asset i as defined in equation (7). Because the distributions of return variances are heavily skewed, the logarithm of the return variance is used in equation (7), which allows to make the distributions more normal.

$$\beta_i^{var} = \frac{\text{cov}[\log(RV_i), \log(RV_{market})]}{\text{var}[\log(RV_{market})]} \quad (7)$$

Initially, the variance beta for each stock or index is calculated over the entire horizon for which options data for the specific underlying are available. Over the entire common time period, the return variance over non-overlapping periods of five trading days is calculated for the underlying and the market portfolio. These figures are then used to determine the variance beta according to equation (7). Even though the variances are computed over a comparatively short period of only five trading days, the results are relatively robust to the use of alternative time horizons. The value-weighted portfolio of all stocks traded at the New York Stock Exchange (NYSE), American Exchange (AMEX) and NASDAQ is used as a proxy for the US market portfolio. Daily returns on this portfolio are obtained from the Center for Research in Security Prices (CRSP). For the three US indices and 29 individual stocks a regression of the mean LVRP over the entire horizon on the respective variance beta obtains the results shown in equation (8).

$$LVRP_i = -0.2424 - 0.1044 \cdot \beta_i^{var} + e_i \quad R^2 = 4.94\% \quad (8) \\ (-4.0327) \quad (-0.8976)$$

t-statistics are based on heteroskedasticity-robust standard errors and shown in parentheses. The goodness of fit is relatively poor with only 4.94%. The slope coefficient is negative but insignificant while the intercept is not substantially

smaller in absolute magnitude than the average risk premium and highly significant. Even though this result does not suggest any relation between market variance and the observed variance risk premium, it is misleading. The relation between the variance beta and the variance risk premium for the three US indices alone is virtually linear with an R^2 of 99.90% while the poor goodness of fit for equation (8) is mainly the result of an apparently non-existent relationship between exposure to market variance and realized variance risk premia for individual stocks. Of course, due to the limited number of only three index observations, this figure is far from meaningful.

In order to increase the number of observations, the regression approach is changed in the following way. For each of the three US indices, the mean logarithmic variance risk premium is calculated on a yearly basis for all calendar years between 1996 and 2014. The yearly mean logarithmic variance risk premium is defined as the average 22 trading-day variance swap return on all variance swaps whose initiation date falls into the respective calendar year. The corresponding variance beta for the specific year is calculated as outlined before using the returns that occurred in a particular calendar year.

Equation (9) shows the result for a pooled OLS regression with cluster-robust standard errors to control for the correlation between observations of the same index in different years.

$$LVRP_{it} = 0.1746 - 0.7018 \cdot \beta_{it}^{var} + e_{it} \quad R^2 = 16.66\% \quad (9)$$

(1.79) (-6.94) $N = 51$

For the three US indices, there is a pronounced relation between the variance beta and the average logarithmic variance risk premium. The slope coefficient is negative and highly significant at the 1%-level whereas the intercept is not significant at the 5% level. These results are in line with the expectation that the average variance risk premium should be more negative for underlyings that have more exposure to overall market variance and suggest that exposure to overall market variance is indeed priced in US stock index options and variance swaps. However, the relatively low R^2 of only 16.66% shows that a significant proportion of overall variability in variance swap returns remains unexplained by the variance beta.

For a similar analysis with the five European stock indices, the Stoxx Europe 600 index is used as a common proxy for a European market portfolio. When the relation between the variance risk premium and the variance beta is examined over the entire common sample, the intercept is negative and highly significant whereas the slope coefficient is positive but insignificant. The positive slope coefficient is counterintuitive since it implies that investors are willing to pay a premium for additional exposure to systematic variance risk, a relation that is not supported by the assumption of risk-averse decision makers who dislike uncertainty. However, due to the limited number of only five observations this result should be considered with the appropriate degree of skepticism. In the corresponding regression with annual observations, the slope coefficient remains positive and insignificant.

In contrast, for the individual stocks, there is no discernible relation between exposure to overall market variance, as measured by the variance beta, and the variance risk premium. Estimating regression (8) and (9) for the stocks only leads to an insignificant slope coefficient and an R^2 indistinguishable from zero but a negative and highly significant intercept. Consequently, there is no evidence that exposure to systematic market variance is priced in variance swaps on individual stocks.

A possible explanation for this observation might be that the observed negative variance risk premium in options on single stocks does not, or only to a small extent, represent compensation for exposure to changes in market variance. There is an increasing amount of literature that documents a common factor structure in the volatility of idiosyncratic returns of stocks, i.e. the return component that cannot be explained by commonly applied factor models (e.g. [Herskovic et al. \(2014\)](#)), and increasing evidence that the risk of common changes in idiosyncratic volatility may indeed be priced (e.g. [Gourier \(2015\)](#), [Cao and Han \(2013\)](#)¹⁰). While standard-asset pricing theory would predict that idiosyncratic risk can be diversified away and should therefore not carry a risk premium, advocates of priced idiosyncratic risk argue that there might be structural issues, such as non-traded assets ([Herskovic et al. \(2014\)](#)), that cause market participants to hold undiversified portfolios. In the context of variance swaps, [Schürhoff and Ziegler \(2011\)](#) find a positive price of idiosyncratic variance risk whereas [Gourier \(2015\)](#) finds that idiosyncratic variance risk carries a negative price. In particular, [Gourier \(2015\)](#) finds that the idiosyncratic component of variance risk explains on average about 80% of the total variance risk premium. As a consequence, the apparently missing relation between exposure to market variance and the variance risk premia in stock options could possibly be explained if the observed premia entirely or in part represented compensation for changes in idiosyncratic variance. In order to test this hypothesis, I follow [Schürhoff and Ziegler \(2011\)](#) and use the excess returns on synthetic variance swaps on the S&P 500 Index as proxy for a systematic variance risk factor or price of overall market variance risk. The excess returns from variance swaps on single stocks and US indices are then regressed on this proxy. Using the variance swap returns on the S&P 500 index as a proxy for a systematic variance factor appears justifiable since the previous regression results show that returns on variance swaps on the three US indices are significantly associated with exposure to overall market variance. Moreover, in a simple factor decomposition [Schürhoff and Ziegler \(2011\)](#) show that index variances are predominantly determined by factor variances and the effect of idiosyncratic returns on index variances is negligible. If a significant proportion of single stock variance swap

¹⁰[Cao and Han \(2013\)](#) examine the returns on delta-hedged options that are sorted according to their idiosyncratic volatilities. They find that the average return from delta-hedging options on stocks with high idiosyncratic volatility is significantly lower than returns from hedging options on stocks with lower idiosyncratic volatility.

returns can then be explained by the systematic variance risk factor, these returns do probably not represent compensation for (common) idiosyncratic variance risk.

Table 3 shows the corresponding regression results. Except for Facebook, the slope coefficient on the systematic variance risk factor is positive and significant at the 1%-level for almost all regressions. R^2 s frequently exceed 30% and achieve values of up to 93.20% and 70.06% for the two indices. This result clearly shows that variance swap returns on different underlyings have a pronounced tendency to move together. The systematic variance risk factor is obviously able to explain a substantial portion of the returns on single stock variance swaps. Even though the alpha is significant for several underlyings, including the two indices, it is not significantly different from zero at the 5%-level for 21 of 29 stocks. This finding supports the view that exposure to the systematic variance risk factor is mostly sufficient to explain variance swap returns and that the previously stated notion of priced idiosyncratic variance risk appears to play only a minor role for explaining single stock variance swap returns, at least for the sample at hand.

5.2.2. Commonly used risk factors

Sources of variance fluctuations and their implications for the pricing of variance risk

The results in the previous section suggest that returns on index as well as on single stock variance swaps have a strong tendency to evolve together and are captured well by the systematic variance risk factor. In this subsection, the objective is to further examine which risk factors precisely drive variance swap returns and to assess whether these returns can be explained by commonly used risk factors or whether exposure to changes in market variance indeed appears to command a separate risk premium. A prerequisite for variance risk to be priced as an independent risk factor is that return variance can evolve independently and affect aggregate consumption. Thus, it is important to consider the possible sources of changes in return variance and their implications for the pricing of variance risk.

Most reduced form option pricing models embed a negative correlation between equity returns and volatility (e.g. Cox (1996), Heston (1993)). Indeed, the presence of this negative correlation, often referred to as asymmetric volatility, is empirically well documented (e.g. Glosten et al. (1993)) and most apparent during significant market downturns that frequently come along with substantial increases in market volatility (Wu (2001)).

The phenomenon of asymmetric volatility is often explained by either the so-called leverage effect or the volatility feedback effect. According to the leverage effect hypothesis, a substantial stock price decline causes the market value of equity to fall more rapidly than the market value of a fixed amount of debt, thus increasing the debt-to-equity ratio. As a consequence the stock's risk rises, which in turn increases volatility (Black (1976)). Since the payoff of a variance swap is positively related to the level of realized vari-

ance, it obviously offers insurance against increases in market or stock volatility and, given the empirically documented negative correlation between market volatility and returns, also insurance against substantial declines in stock prices. This provides a rationale for the negative average return on long variance swaps. The important implication of the leverage effect hypothesis, however, is that volatility does not evolve independently but rather varies as a result of fluctuations in the stock price. If the functional relationship is even deterministic as for example in the model proposed by Dupire (1994), this also has important ramifications in the context of option pricing. An instantaneous volatility that evolves according to a deterministic function of the price of the underlying does not necessarily cause the market to be incomplete and can still allow to hedge an option solely with the underlying (Dupire (1994)). In such a setting, return variance does not represent an independent risk factor that could command a risk premium and the arguments outlined in section 3 that link the variance risk premium to market participants' inability to perfectly hedge an option position would not apply. Consequently, it may well be the case that the returns on variance swap contracts solely result from their directional stock price exposure rather than from explicitly priced variance risk. If this is indeed the case, the classical capital asset pricing model should be able to fully explain variance swap returns. A regression of variance swap returns on the excess market return can therefore reveal whether the observed variance risk premium reflects compensation beyond the correlation between market and variance swap returns.

The volatility feedback effect offers a second potential explanation for the observed negative correlation between volatility and equity returns. In this case, however, the causality between negative equity returns and increasing volatility is reversed as compared to the leverage effect. Here, volatility varies over time, for instance as a result of changes in business risk (Carr and Wu (2011)). Because volatility can fluctuate independently, it can theoretically command a risk premium and is indeed assumed to be priced. If volatility then rises, the required return on equity increases and stock prices fall, thus causing the negative correlation (Wu (2001)).

A third channel through which market returns can negatively interact with return variance is the so-called self-exciting behavior of financial markets which describes the phenomenon that substantial negative financial events seem to increase the likelihood of similar events to follow (Carr and Wu (2011)). Conceptually, this third channel is related to time-varying jump intensity (Carr and Wu (2011)).

Intuitively, all three explanations imply a negative market beta for variance swaps. However, only if volatility can fluctuate independently, it can constitute an additional risk factor that justifies a risk premium.

Capital Asset Pricing Model

In order to test whether the capital asset pricing model can explain observed variance risk premia, variance swap returns are regressed on the market excess return. For US underlyings, daily overlapping regressions as well as non-

Table 3: Regression of logarithmic variance risk premium on return on systematic variance risk factor

Note: Entries report the OLS estimates and t-statistics (in parentheses) of regressions of the continuously compounded excess return on variance swaps over a horizon of 22 trading days on Var_{market} . Var_{market} denotes the systematic variance risk factor or market price of overall market variance risk that is proxied by the continuously compounded excess return over a horizon of 22 trading days on synthetic variance swaps on the S&P 500 Index. t-statistics are adjusted for serial correlation according to the method of Newey and West (1987) with a lag length of 22 days. N denotes the total number of observations.

Underlying	α		Var_{market}		R^2	N
Dow Jones Industrial	0.031	(2.268)	0.961	(53.952)	93.20%	2049
NASDAQ100	0.059	(2.537)	0.747	(24.823)	70.06%	3265
Alcoa	-0.039	(-0.669)	0.463	(6.107)	32.68%	776
Altria (Philip Morris)	-0.169	(-2.719)	0.363	(4.712)	8.23%	2103
Amazon	-0.023	(-0.507)	0.395	(7.385)	13.45%	2211
American Express	0.158	(4.571)	0.695	(16.063)	55.36%	1864
Amgen	-0.066	(-1.638)	0.483	(10.918)	23.31%	2223
Analog Devices	0.056	(0.895)	0.493	(7.536)	36.17%	662
Apple	-0.078	(-1.921)	0.446	(8.047)	22.60%	2282
Bank of America	0.113	(1.879)	0.642	(9.273)	39.68%	1423
Boeing	-0.009	(-0.264)	0.485	(10.315)	33.04%	1966
Cisco	0.043	(1.028)	0.539	(11.479)	27.75%	1850
Exxon Mobil	0.027	(0.684)	0.671	(12.879)	51.36%	1679
Facebook	-0.280	(-2.408)	-0.050	(-0.229)	0.15%	501
General Electric	0.077	(2.462)	0.670	(16.918)	53.68%	1800
Home Depot	0.008	(0.215)	0.614	(13.373)	44.67%	1732
IBM	-0.024	(-0.624)	0.588	(13.531)	33.50%	2770
Johnson & Johnson	-0.067	(-1.375)	0.699	(10.159)	42.43%	1203
McDonald's	-0.182	(-4.333)	0.452	(9.918)	32.39%	1358
Merck	0.008	(0.172)	0.550	(11.732)	25.80%	1963
Metlife	0.153	(2.193)	0.621	(6.990)	51.90%	914
Microsoft	0.015	(0.421)	0.491	(8.943)	28.20%	2318
Monsanto	-0.028	(-0.688)	0.513	(9.936)	42.92%	1240
Nike	0.002	(0.029)	0.548	(8.604)	27.84%	1403
Pfizer	0.039	(0.924)	0.606	(10.807)	34.19%	1398
Procter & Gamble	-0.057	(-0.820)	0.649	(8.760)	34.88%	1555
Starbucks	-0.022	(-0.416)	0.442	(7.149)	23.10%	1496
Tesla	-0.303	(-3.156)	0.024	(0.205)	0.05%	718
Valero	0.031	(0.582)	0.400	(5.760)	20.96%	1674
Verizon	-0.083	(-1.857)	0.441	(7.363)	30.08%	1352
WalMart	-0.081	(-2.157)	0.537	(12.801)	35.18%	1815

overlapping regressions of 22 trading day variance swap returns on the corresponding excess market returns are conducted with the same proxy for the market portfolio as outlined above. The excess market return is calculated as the 22 trading day return on the market portfolio minus the four-week Treasury-Bill rate. For the regressions in Panel A of Table 4, the sample covers only the time period from July 31st 2001 until August 31st 2014, which is due to the fact that the daily time series of four-week Treasury bill rates is available only from July 2001 on. The results reported for regressions in Panel B generally cover the entire sample period for the respective underlying.

For the five European indices, results refer to non-overlapping monthly regressions of 22 trading day variance swap returns on the monthly excess return on a proxy for the Euro-

pean market portfolio. This excess return is publicly available at a monthly frequency from Kenneth French's online data library.¹¹ The regressions are estimated by OLS and standard errors are adjusted for serial correlation according to the method by Newey and West (1987) with 22 lags for the US underlyings in Panel A and unadjusted for the non-overlapping regressions shown in Panel B.

As predicted, the beta is negative and significant at the 1% level for almost all underlyings. The negative beta is

¹¹The internet address is: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#International. The countries considered in the European market portfolio proxy include Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom. This information is also available at Kenneth French's website.

Table 4: CAPM regressions

Note: Entries in Panel A report the OLS estimates and t-statistics (in parentheses) of overlapping regressions of 22-day continuously compounded variance swap returns on the corresponding market excess return which is proxied by the 22-day excess return on the value-weighted portfolio of all stocks traded at the NYSE, NASDAQ, and AMEX. t-statistics are computed based on standard errors that are adjusted according to the method of Newey and West (1987) with a lag length of 22 days. Entries in Panel B report the OLS estimates and t-statistics (in parentheses) of non-overlapping regressions of 22-day continuously compounded variance swap returns on the corresponding market excess return. For the US underlying, the excess market return is proxied by the monthly excess return on the value-weighted portfolio of all stocks traded at the NYSE, NASDAQ, and AMEX. For the European indices, the market excess return refers to the excess return on a proxy for a European market portfolio taken from Kenneth French's website. t-statistics are not adjusted for serial correlation. N denotes the total number of observations.

Underlying	Panel A:						Panel B:					
	alpha	t-stat	beta	t-stat	R ²	N	alpha	t-stat	beta	t-stat	R ²	N
S&P 500 Index	-0.499	-15.476	-7.001	-7.393	37.30%	2587	-0.482	-14.006	-7.162	-8.177	31.60%	213
Dow Jones Industrial	-0.444	-12.835	-7.140	-6.997	36.60%	2056	-0.419	-10.012	-8.817	-7.189	39.10%	134
NASDAQ100	-0.357	-12.621	-5.863	-6.955	33.40%	2387	-0.305	-9.574	-6.254	-7.873	29.60%	207
Alcoa	-0.264	-6.189	-6.475	-8.655	39.80%	694	-0.237	-3.610	-6.953	-4.299	35.90%	50
Altria (Philip Morris)	-0.416	-6.228	-3.558	-2.971	5.60%	1464	-0.369	-6.660	-2.761	-2.391	3.60%	155
Amazon	-0.273	-6.259	-3.563	-4.869	8.00%	1978	-0.180	-3.362	-2.785	-2.485	3.60%	154
American Express	-0.198	-5.019	-5.436	-6.680	25.00%	1470	-0.118	-2.465	-4.356	-4.278	14.50%	129
Amgen	-0.332	-7.230	-4.255	-6.073	14.00%	1827	-0.263	-5.776	-3.587	-3.729	9.50%	154
Analog Devices	-0.278	-5.224	-2.790	-2.945	9.80%	467	-0.111	-1.497	-2.498	-1.974	6.50%	46
Apple	-0.319	-8.883	-4.396	-5.748	17.40%	1901	-0.260	-5.552	-5.473	-4.433	17.00%	153
Bank of America	-0.202	-3.667	-6.416	-6.601	30.20%	1094	-0.108	-1.868	-6.120	-4.842	25.70%	89
Boeing	-0.268	-7.969	-4.090	-6.849	19.70%	1721	-0.199	-4.800	-4.341	-5.430	17.00%	134
Cisco	-0.281	-5.135	-5.036	-5.136	18.20%	1135	-0.211	-4.220	-5.856	-6.476	21.40%	127
Exxon Mobil	-0.285	-7.204	-6.239	-6.582	33.30%	1573	-0.287	-5.845	-6.365	-4.865	28.40%	114
Facebook	-0.158	-1.728	-5.764	-2.180	6.00%	498	-0.265	-3.430	-1.911	-0.367	0.80%	24
General Electric	-0.251	-5.762	-5.241	-7.508	25.20%	1248	-0.205	-4.339	-3.801	-3.880	13.10%	119
Home Depot	-0.354	-9.578	-4.321	-7.794	19.60%	1429	-0.224	-4.232	-4.796	-5.474	16.80%	114
IBM	-0.394	-10.076	-4.495	-6.649	17.30%	1980	-0.316	-8.291	-4.685	-6.646	15.80%	190
Johnson & Johnson	-0.410	-7.405	-6.131	-5.423	26.60%	959	-0.442	-6.620	-3.496	-2.588	8.40%	82
McDonald's	-0.403	-10.156	-3.947	-4.251	19.40%	1291	-0.376	-7.209	-4.276	-3.254	16.90%	92
Merck	-0.321	-6.139	-4.241	-4.705	10.90%	1458	-0.255	-4.473	-3.510	-3.217	6.60%	131
Metlife	-0.160	-3.045	-5.765	-4.044	25.10%	943	-0.109	-1.426	-7.518	-3.034	29.30%	64
Microsoft	-0.260	-6.102	-4.088	-4.570	17.40%	1508	-0.200	-4.860	-3.910	-4.658	13.00%	160
Monsanto	-0.264	-7.202	-4.361	-5.492	23.90%	1284	-0.216	-4.409	-4.971	-4.186	21.80%	90
Nike	-0.293	-5.942	-5.222	-4.880	21.00%	1149	-0.226	-3.380	-4.502	-3.076	12.40%	87
Pfizer	-0.328	-5.584	-5.091	-5.498	18.90%	912	-0.237	-4.311	-2.937	-2.927	7.80%	94
Procter & Gamble	-0.463	-9.190	-4.592	-4.744	20.40%	1302	-0.363	-5.483	-4.735	-3.354	13.10%	97
Starbucks	-0.269	-6.179	-4.384	-5.491	15.40%	1287	-0.240	-4.497	-4.323	-3.854	13.30%	103
Tesla	-0.319	-4.777	-0.004	-0.003	0.00%	728	-0.393	-4.299	0.249	0.108	0.00%	45
Valero	-0.152	-3.944	-4.185	-5.471	15.90%	1741	-0.149	-2.938	-4.398	-3.504	11.50%	121
Verizon	-0.298	-8.135	-4.086	-5.130	20.50%	1352	-0.258	-4.686	-4.703	-3.822	20.40%	85
Walmart	-0.405	-11.140	-4.779	-8.185	26.50%	1462	-0.310	-6.541	-4.396	-5.186	16.10%	125
AEX Index	-	-	-	-	-	-	-0.258	-3.306	-9.058	-6.258	51.60%	26
CAC40 Index	-	-	-	-	-	-	-0.236	-3.778	-8.529	-7.756	61.10%	38
DAX Index	-	-	-	-	-	-	-0.311	-5.679	-7.061	-7.245	49.10%	63
Euro Stoxx 50	-	-	-	-	-	-	-0.267	-3.823	-8.460	-6.391	51.60%	30
SMI Index	-	-	-	-	-	-	-0.419	-3.613	-6.745	-3.341	23.40%	22

consistent with the empirically found negative correlation between return variance and equity returns. However, with a R^2 somewhere between 10% and 30% for the majority of regressions a substantial portion of the variation in variance swap returns remains unexplained. Furthermore, with only few exceptions, the alpha is negative and highly significant at the 1%-level, with the absolute magnitude of the alpha not significantly lower than the average variance risk premium reported in Table 2. Indeed, based on the average 22 trading day market excess return for the US market of approximately 0.89% over the sample period in Panel A, the market beta accounts on average for only approximately 7.01%¹² of the logarithmic variance risk premium across all US underlyings. A comparison of the two panels moreover shows that the results are qualitatively unaffected by the sample frequency.

For the five European indices, the situation is very similar. The market beta as well as the alpha is negative and highly significant leading to the same conclusion as for the US underlyings. What is striking, however, is that the R^2 s for the European indices are mostly in the range of 50% and therefore substantially higher than for the US indices, which suggests that the selected market portfolio proxy is appropriate and captures the dynamic in European variance swap returns quite well.

All in all, the fact that the market beta cannot fully account for the realized variance swap returns represents an indication that the observed variance risk premium reflects compensation for other factors than the correlation between variance swap returns and market excess returns.

Fama-French three-factor model

Considering the fact that the CAPM has been found to be insufficient to explain stock returns (Fama and French (2004)), there is no apparent reason to believe the CAPM should perform much better in explaining variance swap returns. Thus, in order to evaluate whether these returns can be better explained when additional factors are considered, they are regressed on the excess market return and the two additional factors identified by Fama and French (1993), a firm size factor (SMB) and a book-to-market factor (HML). The excess return realizations for the market portfolio proxy and the SMB and HML factor portfolios, both for the US and Europe, are publicly available at Kenneth French's website at a monthly frequency.¹³ The market portfolio proxies are the same as those used for the CAPM regressions. Since the excess returns are published on a monthly basis, the regression results in Table 5 refer to monthly non-overlapping regressions.

The magnitude of neither the alpha nor the market beta changes substantially compared to the values obtained in the CAPM regressions. Moreover, both parameters remain highly significant. Consistent with the findings of Carr and

Wu (2009), the loading on the SMB factor is usually negative and often significant for the US underlyings, which suggests a negative correlation not only between return variance and market returns but also between return variance and returns on the SMB factor. Long variance swaps thus appear to offer insurance against a situation in which small stocks underperform relative to large stocks. The sign of the loadings on the HML factor is less persistent and the loading is mostly insignificant. For the European indices the situation is again very similar. Altogether, the persistently negative regression alphas suggest that variance swap returns reflect compensation beyond the swap's pure correlation with the market portfolio and that commonly used risk factors are not able to explain this additional compensation. The results therefore point towards additional risk factors that command a persistent and economically substantial risk premium.

5.2.3. Jump and variance risk

Construction of risk-factor-mimicking portfolios

Thus far, commonly used risk factors do not appear to be able to satisfactorily explain variance swap returns, which is indeed supportive of the notion that these returns reflect compensation for one or more additional priced risk factors. As a consequence, I construct alternative factor portfolios that allow to gain separate exposure to the risks that theory suggests to be priced in option returns and evaluate whether the returns on these risk factor portfolios help to explain variance swap returns.

The total realized quadratic variation is the result of continuous as well as of discontinuous price movements and the uncertainty associated with each of these kinds of movements may induce investors to command separate risk premia for them. Indeed, as outlined in section 3 and noted by Bollerslev and Todorov (2011), the commonly observed variance risk premium does not only reflect compensation for stochastic volatilities but also for time-varying jump intensities and fears for jump tail events. The latter can explain up to three quarters of the variance risk premium (Bollerslev and Todorov (2011)). Especially the fact that jumps are typically required in models to generate a variance risk premium whose magnitude is reconcilable with that observed in reality (e.g. Drechsler (2013)) or in option pricing models to fit certain characteristics of the cross-section of option prices (e.g. Pan (2002)) may point to the importance of jump risk when explaining the variance risk premium.

Thus, to examine the degree to which the returns on variance swaps represent compensation for presumed jump risk, i.e. risk associated with discontinuous price movements, and diffusive risk, I construct two factor portfolios from traded options that are intended to offer exposure to one of the two risk factors while being relatively unaffected by the other. For this purpose, I follow Cremers et al. (2015) who use the same factor portfolios to examine the influence of jump and volatility risk on the cross-section of stock returns.

On any given day t in the sample, I construct two delta-neutral straddles with different maturities. For the US un-

¹²For each underlyings, this contribution is calculated as $marketbeta \cdot 0.89\% / (\alpha + marketbeta \cdot 0.89\%)$.

¹³The internet address is: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Table 5: Fama-French Three-Factor regressions

Note: Entries report the OLS-estimates and t-statistics (in parentheses) of monthly, non-overlapping regressions of 22-day continuously compounded variance swap returns on the corresponding excess return on the market portfolio and the two additional factors, SMB and HML, identified by Fama and French (1993). For the US, the market excess return is proxied by the 22-day excess return on the value-weighted portfolio of all stocks traded at the NYSE, NASDAQ, and AMEX. For the European indices, the market excess return refers to the excess return on a proxy for a European market portfolio retrieved from Kenneth French's website. t-statistics are not adjusted for serial correlation. N denotes the total number of observations.

Underlying	alpha	t-stat	market	t-stat	SMB	t-stat	HML	t-stat	R ²	N
S&P 500 Index	-0.474	-13.818	-7.020	-7.642	-2.073	-1.585	-1.780	-1.578	33.20%	213
Dow Jones Industrial	-0.423	-10.521	-9.717	-7.252	1.164	0.636	3.798	1.840	41.10%	134
NASDAQ100	-0.299	-9.323	-5.848	-7.217	-2.552	-2.893	-0.832	-0.808	31.90%	207
Alcoa	-0.268	-4.114	-6.933	-4.305	-2.553	-1.530	4.212	2.364	44.10%	50
Altria (Philip Morris)	-0.345	-5.994	-2.871	-2.433	-2.698	-1.913	-3.409	-2.176	6.40%	155
Amazon	-0.174	-3.309	-2.160	-1.831	-2.967	-1.860	-4.557	-2.476	7.90%	154
American Express	-0.117	-2.460	-4.296	-4.222	-0.978	-0.715	-1.495	-1.060	15.30%	129
Amgen	-0.253	-5.641	-3.408	-3.362	-1.988	-1.240	-2.081	-1.285	11.40%	154
Analog Devices	-0.067	-0.904	-2.043	-1.629	-3.645	-1.987	-2.790	-1.696	15.00%	46
Apple	-0.261	-5.877	-5.282	-4.240	-0.680	-0.573	0.291	0.169	17.20%	153
Bank of America	-0.093	-1.697	-6.328	-5.769	-2.642	-2.138	-4.000	-2.744	32.30%	89
Boeing	-0.202	-4.903	-3.753	-4.321	-2.962	-1.744	-0.366	-0.271	18.90%	134
Cisco	-0.216	-4.439	-5.654	-6.032	-0.140	-0.105	1.332	1.081	22.10%	127
Exxon Mobil	-0.286	-5.783	-6.261	-4.543	-0.581	-0.334	-0.055	-0.030	28.50%	114
Facebook	-0.287	-3.932	-2.697	-0.561	2.281	0.443	6.983	1.823	4.90%	24
General Electric	-0.202	-4.350	-3.885	-4.134	-2.034	-1.553	-2.115	-1.875	16.00%	119
Home Depot	-0.220	-4.059	-4.465	-4.996	-1.242	-0.736	0.261	0.221	17.60%	114
IBM	-0.308	-8.118	-4.378	-6.029	-2.629	-2.231	-1.391	-1.148	18.30%	190
Johnson & Johnson	-0.421	-6.001	-3.530	-2.369	-2.622	-1.162	-2.520	-1.102	11.60%	82
McDonald's	-0.379	-7.339	-4.638	-2.927	1.083	0.677	1.632	0.601	17.90%	92
Merck	-0.257	-4.474	-3.399	-2.827	-0.123	-0.064	0.575	0.344	6.70%	131
Metlife	-0.107	-1.399	-8.161	-2.621	0.820	0.244	3.120	0.799	31.00%	64
Microsoft	-0.192	-4.874	-3.594	-4.154	-2.993	-2.803	-1.209	-0.988	16.50%	160
Monsanto	-0.216	-4.421	-4.154	-3.008	-2.919	-1.337	-0.847	-0.356	23.40%	90
Nike	-0.244	-3.734	-3.616	-2.493	-5.066	-2.219	-2.298	-1.161	16.50%	87
Pfizer	-0.230	-4.355	-2.867	-2.946	-2.519	-2.051	-2.429	-1.834	11.80%	94
Procter & Gamble	-0.357	-5.439	-4.066	-2.583	-4.183	-1.158	0.570	0.321	18.90%	97
Starbucks	-0.238	-4.362	-4.597	-4.080	0.325	0.249	-2.208	-1.465	15.30%	103
Tesla	-0.373	-3.771	-0.975	-0.344	4.394	0.802	-5.421	-0.874	4.80%	45
Valero	-0.150	-3.052	-5.179	-3.570	2.468	0.971	1.470	0.680	12.60%	121
Verizon	-0.251	-4.420	-4.725	-3.470	-1.505	-0.678	1.911	0.981	21.80%	85
WalMart	-0.304	-6.276	-4.369	-4.858	-0.952	-0.475	-0.996	-0.703	16.60%	125
AEX Index	-0.246	-4.060	-8.005	-4.824	0.083	2.233	-0.021	-0.540	58.60%	26
CAC40 Index	-0.237	-4.577	-8.238	-6.063	0.045	1.426	-0.004	-0.142	63.00%	38
DAX Index	-0.287	-5.337	-8.039	-7.546	0.038	1.160	0.033	0.997	50.70%	63
Euro Stoxx 50	-0.251	-4.033	-8.167	-5.347	0.056	1.529	0.000	-0.014	54.80%	30
SMI Index	-0.396	-4.317	-5.731	-2.400	0.091	1.628	-0.025	-0.451	31.40%	22

derlyings, these straddles are constructed from options on the S&P 500 index. The choice falls on this specific index because the corresponding options are highly liquid, have the highest number of valid price quotes of all underlyings and cover the longest time period in the sample. Moreover, since the S&P 500 represents the most diversified portfolio among the three US indices considered here, straddles constructed from options on it are most likely to capture exposure to purely systematic risks. In case of the European indices, separate straddle portfolios are constructed for each of the indices from options on the respective index. In general, the same criteria apply as those outlined in section 4.2 with the exception that in-the-money options are not excluded since the

construction of straddles involves put and call options with identical strike and time to maturity that are close to being at-the-money.

Following [Cremers et al. \(2015\)](#), the straddle with the shorter maturity is constructed from the call and put options that are closest to being at-the-money and which expire in the calendar month following the month during which the straddle is constructed while the at-the-money options used to form the straddle with the longer maturity expire in the calendar month that follows. Each straddle return is constructed by solving the following problem

$$w_{t-1}^c \delta_{t-1}^c + (1 - w_{t-1}^c) \delta_{t-1}^p = 0, \quad (10)$$

$$R_t^S = w_{t-1}^c R_t^C + (1 - w_{t-1}^c) R_t^P, \quad (11)$$

where w_t^c denotes the weight of the call in the straddle, R_t^C and R_t^P denote the one-day returns on the call and put, respectively, δ_t^c and δ_t^p denote the Black and Scholes (1973) deltas relative to the option price of the call and put and R_t^S is the weighted one-day return on the straddle, all at time t . In a second step, the JUMP factor which is intended to capture jump risk, and the VOL factor which is intended to capture the diffusive component of volatility risk, are constructed as a weighted combination of the previously created delta-neutral straddles with different maturities.

In order to construct the JUMP factor, a long position in one short-maturity straddle is combined with a short position in n long-maturity straddles such that the combined position is vega-neutral and thus insensitive to moderate changes in implied volatility. Analog to Fournier and Jacobs (2015), the one-day return on the JUMP factor, $JUMP_t$, is thus given by

$$JUMP_t = R_t^{S1} - \frac{Vega_{t-1}^{S1}}{Vega_{t-1}^{S2}} R_t^{S2}, \quad (12)$$

where R_t^{S1} and R_t^{S2} denote the one-day returns on the short- and long-maturity straddle, respectively and $Vega_t^{S1}$ and $Vega_t^{S2}$ denote the Black and Scholes (1973) vegas of the short- and long-maturity straddles relative to their market prices, all at time t . $JUMP_t$ can also be interpreted as the return on a portfolio that consists of a long position in the short-maturity straddle worth \$1 and a short position in the long-maturity straddle worth $Vega_{(t-1)}^{S1}/Vega_{(t-1)}^{S2}$ dollars. Since the vega of an option increases with time to maturity, the number of long-maturity straddles, n , is always lower than the number of short-maturity straddles held in the combination. At the same time, the gamma of an at-the-money option is higher when the time to expiration is lower. Thus, the combined position of the two delta-neutral straddles has positive gamma. In particular this positive gamma constitutes the theoretical exposure to jumps. Whereas the option delta postulates a linear relation between the option price and the price of the underlying, the option gamma measures the convexity in this relation. Due to this convexity, the established hedge is not perfect (Hull (2009)) and the corresponding hedging error increases with the magnitude of changes in the price of the underlying. Moreover, this hedging error always works in favor of a long position in the constructed portfolio since, holding everything else constant, a given increase in the price of the underlying causes a stronger price appreciation in the short-term straddle (that is held long) than in the long-term straddle (that is held short). Similarly, a given reduction in the price of the underlying causes a smaller price depreciation in the short-term straddle than in the long-term straddle. At the same time, however, the theta of an at-the-money option is usually more negative when the remaining time to expiration is shorter (Hull (2009)). Because the theta of the short-term straddle is more negative than that of the long-term straddle and the number of short-term straddles held is larger, the overall combination has a negative theta. Thus,

holding everything else constant, the option portfolio will lose value from one day to the next. Consequently, the overall return on the jump factor will usually be negative and positive only when the hedging error overcompensates for the value reduction that results from the time decay. This will usually be the case when comparatively large movements in the price of the underlying, i.e. jumps, occur.

The VOL factor, in contrast, is specifically designed to be relatively unexposed to jump risk and to only capture the diffusive component of volatility risk. In order to achieve this goal, it is constructed to be delta-neutral and gamma-neutral but to have positive vega. Thus, the VOL factor consists of a long position in one long-term straddle and a short position in n short-term straddles so that the one-day return on the VOL factor, VOL_t , is given by

$$VOL_t = R_t^{S2} - \frac{Gamma_{t-1}^{S2}}{Gamma_{t-1}^{S1}} R_t^{S1}, \quad (13)$$

where $Gamma_t^{S1}$ and $Gamma_t^{S2}$ denote the time- t Black and Scholes (1973) gammas of the short- and long-maturity straddles relative to their prices.

Empirical characteristics of risk-factor-mimicking portfolios: Average returns

Table 6 shows descriptive statistics for the JUMP and VOL factor portfolios for both discrete and continuous returns. In case of the factor portfolios for the US, the JUMP factor earns substantial negative average returns for both forms, discrete and continuous, even though they are significantly negative only for the continuous form. Note, however, that the insignificance for the discrete form of the JUMP factor is due only to two extreme outliers of +83.68% and +45.93%. Excluding these returns leads to a t-statistic of -2.096.

The picture is similar for the VOL factor which also earns significantly negative average returns although the magnitude is somehow smaller than for the JUMP factor. The low correlation between the two factors of only -0.14 moreover suggests that they capture mostly unrelated return sources. Nevertheless, the negative correlation between the two factors is somewhat surprising since for example Todorov (2010) finds that price jumps are associated with corresponding spikes in stochastic volatility.

Figure 2 shows the cumulative returns on the JUMP and VOL factors over the period from January 5th 1996 to August 29th 2014 and further illustrates the economic magnitude of the negative average returns investors would have incurred if they had revolved the positions daily.

Even though long positions in both factors would have earned substantial positive returns around some major events such as the Asian crisis, or the bankruptcy of Lehman Brothers, the cumulative returns quickly converge to -100%. Considering the fact that both factor are market-neutral per construction, these results suggest that the JUMP and VOL factors indeed capture different factors that are both heavily priced by the market, presumably jump risk and volatility risk.

Table 6: Descriptive statistics for the risk factor mimicking portfolios JUMP and VOL

Note: JUMP denotes the daily return on a delta-neutral, vega-neutral and positive gamma calendar spread option strategy. VOL denotes the daily return on a delta-neutral, gamma-neutral and positive vega calendar spread option strategy. The additives (disc.) and (cont.) indicate the use of discrete and continuously compounded daily returns, respectively. t-statistics are shown in parentheses. N denotes the number of observations and SD denotes the standard deviation.

			Annualized Mean		Annualized SD	Annualized Sharpe Ratio	Daily Median	Skewness	Kurtosis	N
S&P 500	JUMP	(disc.)	-0.172	(-1.347)	0.544	-0.317	-0.005	5.157	95.821	4597
		(cont.)	-4.856	(-51.766)	0.398	-12.191	-0.016	-1.759	25.465	4597
	VOL	(disc.)	-0.138	(-2.332)	0.252	-0.549	-0.001	2.376	46.496	4597
		(cont.)	-0.122	(-2.011)	0.258	-0.474	-0.001	0.731	35.327	4597
AEX	JUMP	(disc.)	0.213	(0.877)	0.496	0.431	-0.004	3.322	39.396	1057
		(cont.)	-4.318	(-24.836)	0.354	-12.193	-0.014	0.583	32.682	1057
	VOL	(disc.)	0.185	(1.536)	0.245	0.754	-0.001	5.804	80.325	1057
		(cont.)	0.158	(1.418)	0.226	0.696	0.000	4.529	57.356	1057
CAC40	JUMP	(disc.)	0.591	(2.231)	0.478	1.236	-0.003	1.862	8.917	829
		(cont.)	-5.136	(-26.264)	0.353	-14.558	-0.017	-1.516	8.755	829
	VOL	(disc.)	-0.052	(-0.477)	0.196	-0.264	-0.001	0.564	6.292	829
		(cont.)	0.014	(0.135)	0.193	0.075	0.000	0.720	8.169	829
DAX	JUMP	(disc.)	0.016	(0.100)	0.480	0.033	-0.004	4.836	60.087	2306
		(cont.)	-4.576	(-37.536)	0.367	-12.480	-0.015	-0.789	112.567	2306
	VOL	(disc.)	0.038	(0.627)	0.182	0.208	0.000	2.111	37.169	2306
		(cont.)	0.102	(1.361)	0.225	0.453	0.000	11.306	361.090	2306
Eurostoxx	JUMP	(disc.)	0.229	(1.003)	0.485	0.471	-0.005	2.947	28.961	1153
		(cont.)	-4.497	(-27.107)	0.353	-12.742	-0.015	-0.347	33.886	1153
	VOL	(disc.)	-0.079	(-0.806)	0.210	-0.379	-0.001	-1.068	49.187	1153
		(cont.)	-0.041	(-0.446)	0.197	-0.210	-0.001	0.604	33.659	1153
SMI	JUMP	(disc.)	1.180	(2.277)	0.650	1.816	-0.003	3.712	30.644	400
		(cont.)	-5.868	(-15.198)	0.484	-12.120	-0.019	-3.473	34.358	400
	VOL	(disc.)	0.118	(0.572)	0.260	0.456	-0.001	0.464	8.787	400
		(cont.)	0.024	(0.115)	0.260	0.092	0.000	-1.157	19.515	400

Moreover, the negative sign of the average returns is consistent with models where investors dislike jumps or fear a deteriorating investment opportunity set due to increased variance and are therefore willing to accept lower or even negative average returns on assets or combinations of such assets that offer a hedge against jumps or increased variance (e.g. [Drechsler \(2013\)](#)).

For the five European indices, the picture is remarkably different. Discrete returns on the JUMP factor are on average positive for all indices even though insignificant for three of them. However, for the CAC40 and the SMI, the annualized returns are substantial and significant. In general, positive average returns on the JUMP and VOL factors are counter-intuitive and contradict the idea that investors are willing to accept lower average returns on assets that offer insurance against undesirable states of the world. In this context, a possible bias that results from the factor portfolio construction methodology should be taken into account. This bias could particularly concern the risk factor returns for the CAC40 and SMI. As can be seen in Table 6, the two indices for which average factor portfolio returns are remarkably high (CAC40 and SMI) are also the two indices with the lowest number of valid days on which straddle returns could be computed. This situation occurs even though the raw data on the AEX

and SMI cover exactly the same time horizon and the data on the CAC40 cover an even longer time period than on the Euro Stoxx 50. The reason for the different numbers of valid days is that the straddles that are used for the two factor portfolios can only be formed when the same ATM-options have positive trading volume on two consecutive trading days, which apparently is not always the case. If one now suspects that investors' trading in options is, at least in part, driven by the desire to insure against substantial negative price movements or rising volatility or otherwise take speculative positions, investors are probably more likely to trade immediately around the time of certain events, for instance when the underlying has incurred unexpected price fluctuations or volatility has unexpectedly risen. Indeed, [Cao and Ou-Yang \(2009\)](#) conclude that differences of opinion are more likely to exist when a big rare event occurs and that option trading should be clustered during the time of such events. However, if this is indeed the case, the likelihood to observe options with positive trading volume on two consecutive trading days is also higher at such times. The potential bias results from the fact that the JUMP and VOL factors are constructed to earn positive returns precisely when such events occur. As a consequence, the likelihood to obtain valid return observations for the JUMP and VOL factors is probably

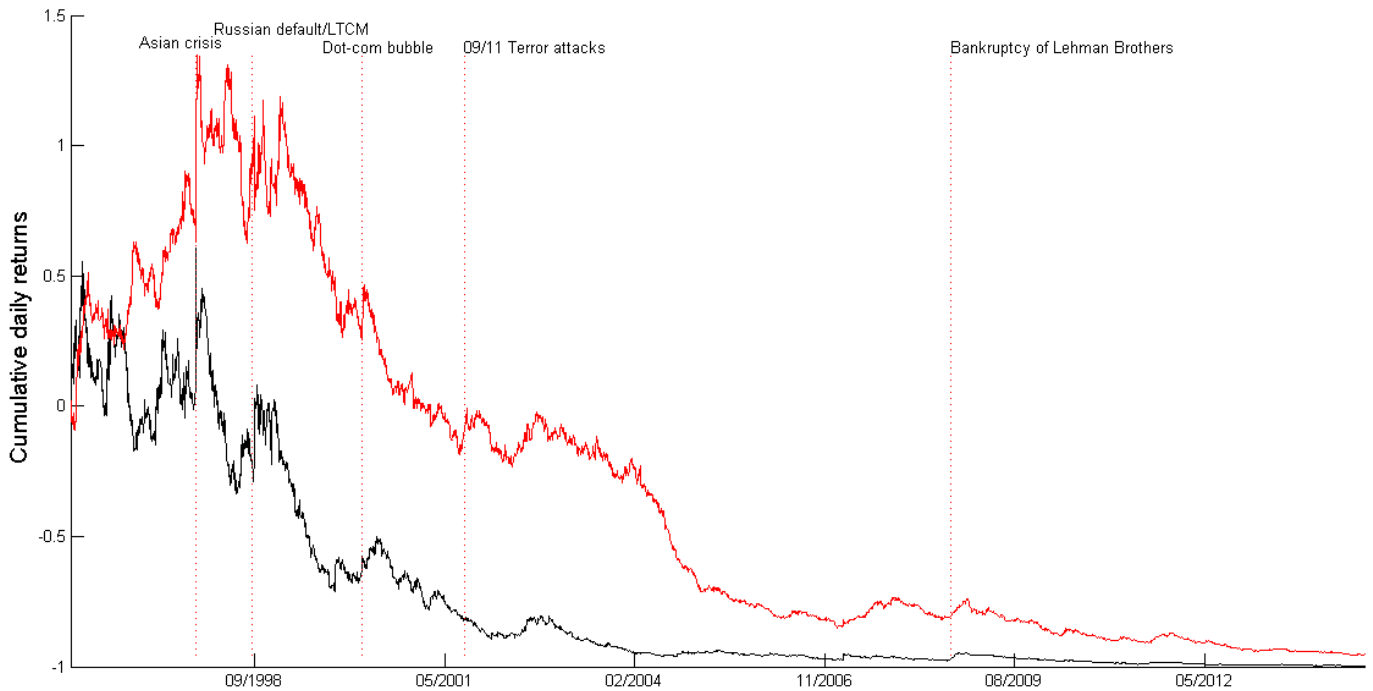


Figure 2: Cumulative daily returns on risk-factor-mimicking portfolios JUMP and VOL

Note: The black line shows the cumulative daily returns on the jump-risk mimicking portfolio JUMP (calculated from discrete returns) over the sample period from January 5th 1996 August 29th 2014. The red line shows the cumulative daily returns on the VOL factor over the same period. Vertical dashed lines indicate major events.

higher on days on which the two factors earn positive returns, which erroneously biases the average return upwards. Since positive returns on the constructed factor portfolios are frequently double-digit, this bias can potentially be substantial. This may partly explain the remarkably high positive average returns on the CAC40 and the SMI reported in Table 6. Note that for the two factors constructed from S&P 500 options, valid returns could be computed on more than 97.8% of all days in the sample so that the previously explained situation is unlikely to bias average returns on these two factors. For the VOL factors, the sign is not persistent across the five European indices and the mean return is never significantly different from zero.

Due to the frequent and often extended time gaps between valid factor return observations for the European indices, cumulative returns are significantly less meaningful than for the US factors. Nevertheless, for the sake of completeness, Figure A1 in the Online-Appendix shows the cumulative returns for the European indices.

As with the variance risk premium before, the returns on the two factor portfolios appear to stem from a non-normal distribution with substantial skewness and kurtosis that can partly be alleviated by using continuous returns instead of discrete. However, as noted by Coval and Shumway (2001) and indicated by the descriptive statistics for the JUMP factor shown above, logarithmic option returns are significantly lower than the raw option returns and this relocation is markedly more pronounced than for variance swap returns.

Due to this, discrete returns will be used for the following analyses.

Empirical characteristics of risk-factor-mimicking portfolios: Determination of jump risk exposure

Because the validity of any conclusion drawn about the extent to which the variance risk premium reflects compensation for jump risk requires that the JUMP factor be able to measure jump risk accurately and not simply other priced factors, I examine whether positive returns on the JUMP factor coincide with jumps indicated by the non-parametric jump detection test of Lee and Mykland (2008) which allows to determine if a jump has occurred at a particular point of time. The underlying idea is that large price movements and associated returns can either be due to a discontinuous movement, i.e. represent a jump in the true sense, or simply be the result of a comparatively high spot volatility at the particular point of time. In order to reliably distinguish the two cases, the authors propose to standardize the return that is to be analyzed by the instantaneous volatility and derive the following test statistic, which determines whether a jump occurred from time $(t - 1)$ to t .

$$\Lambda(t) \equiv \frac{\log(S_t/S_{t-1})}{\widehat{\sigma}(t)}, \quad (14)$$

where $\widehat{\sigma}(t)$ represents the square root of the estimated bipower variation based on the previous W observations of

the spot price S_t at time t .

$$\widehat{\sigma(t)^2} \equiv \frac{1}{W-2} \sum_{j=t-W+2}^{t-1} (|\log(S_j/S_{j-1})| + |\log(S_{j-1}/S_{j-2})|) \quad (15)$$

In constructing the test statistic and computing the realized bipower variation, I follow the recommendation of Lee and Mykland (2008) to use an optimal window size of $W = 16$ for daily observations. The test statistic follows approximately a normal distribution if there is no jump at the testing time, but becomes very large otherwise. A reasonable rejection region can be determined by comparing the test statistic with the usual region of maximums of the test statistic. If the test statistic is outside this usual region, it is probable that the examined price movement represents a jump in the true sense. Thus, the null hypothesis of no jump at a particular point of time can be rejected if $\frac{|\Lambda(t)| - C_n}{S_n}$ exceeds the critical value β^* , where $\beta^* = -\log(-\log(1 - \alpha))$, $C_n = \frac{(2 \log(n))^{1/2}}{c} - \frac{\log(\pi) + \log(\log(n))}{2c(2 \log(n))^{1/2}}$, $S_n = \frac{1}{c(2 \log(n))^{1/2}}$ and $c = \sqrt{2/\pi}$ for a given confidence level α and sample size n .

The analysis whether the JUMP factor truly captures exposure to price jumps is limited to the JUMP factor based on S&P 500 options. The reason is that the JUMP time series for the S&P 500 index is by far the longest and can therefore be expected to offer the most meaningful results.

Figure 3 shows the time series of daily returns on the JUMP factor. The vertical dashed lines indicate the occurrence of jumps determined by the applied jump detection test. Based on a confidence level of $\alpha = 5\%$, the jump test identifies 66 days in the time series of the S&P 500 index on which jumps occur. As can be clearly seen in the figure, jumps in the index are usually accompanied by positive returns on the JUMP factor. In particular, the returns on the JUMP factor are positive for 62 of the 66 days on which a jump is detected with an average daily return of +8.39% on jump days compared to an average daily return of -0.19% on days without a jump. Even though this is a first indication that the JUMP factor indeed captures the effect of price jumps, it is also apparent in Figure 3 that returns on the JUMP factor are positive on many days on which the test does not detect a jump. In this context note that the asymptotic arguments used for setting up the test statistic and the rejection region strongly depend on the assumption that the time period between consecutive observations converges to 0. As a consequence, the use of daily data naturally inhibits the detection of jumps. Lee and Mykland (2008) point out that only 2% of jumps can be detected using daily return data. Thus, the test results should only be regarded as an additional indication to support the theoretical argument for the JUMP factor's exposure to price jumps given above, and not as an independent attempt to prove this exposure.

To further link the returns on the JUMP factor to price jumps, daily JUMP factor returns that occur on days that are no jump days but directly follow such days are compared with the corresponding returns on days that neither immediately

follow nor are jump days. The average daily return on days that immediately follow a jump day is -0.87% compared to -0.18% on other days. A regression of the daily JUMP factor returns on non-jump-days on an indicator variable that is one when the return occurs on a day that immediately follows a jump day but is not a jump day itself, and zero when the return occurs on a day that does not follow a jump day and is not a jump day itself, obtains a t-statistic of -1.8595 for the indicator variable. Thus, it seems that investors require a higher compensation for bearing the risk to which the JUMP factor offers exposure, presumably jump risk, immediately after a price jump has occurred. Conceptually, this observation is plausible with regard to an investor who reevaluates the jump probability and deems the occurrence of a further jump more likely after a jump has occurred. Such a behavior would also be consistent with the empirical evidence that financial markets exhibit strong self-excitation over time (e.g. Ait-Sahalia et al. (2015a)), i.e. that jumps are often followed by further jumps or substantial price movements. Possible reasons that could explain this pattern include credit and liquidity shocks and margin calls that are caused by preceding jump events and then trigger further jumps (Ait-Sahalia et al. (2015a)). An alternative explanation could be investors' increased desire to protect against further market drops as in the model by Bates (2008) where wealth redistributions due to preceding market drops increase average crash aversion.

All in all, the return behavior appears to be reconcilable with the notion that the JUMP factor indeed captures exposure to jump risk, which is therefore assumed for the remainder.

Compensation for jump risk and diffusive risk in variance swap returns

For the following analyses, I follow Cremers et al. (2015) and treat the returns on the two risk factors directly as excess returns that can be used in the regressions. In order to compare variance swap returns and returns on the JUMP and VOL factors over the same time horizon, the cumulative factor returns over 22 trading days are used in the following regressions. Note that the requirement to have returns on 22 consecutive trading days substantially reduces the numbers of available observations compared to those reported in Table 6. For example for the SMI there is no single period with valid returns over 22 consecutive trading days.

Table 7 shows the results of overlapping regressions of the LVRP on the JUMP and VOL factors.

The obtained R^2 s are generally high and exceed 60% for five of the seven indices with more than 90% for the CAC40. As expected, the coefficient on the VOL factor is positive and significant at the 1% level for most regressions, which confirms that exposure to changes in systematic variance is heavily priced not only in variance swaps on indices but also in such on individual stocks. Moreover, consistent with prior studies which find that a substantial portion of the observed variance risk premium represents compensation for jump risk (e.g. Todorov (2010), Bollerslev and Todorov (2011)), the JUMP factor is significant at the 1% level for all underlyings




Figure 3: Daily returns on jump risk proxy

Note: The figure shows the time series of daily discrete returns on the jump risk factor JUMP over the period from January 5th 1996 August 29th 2014. Vertical dashed lines indicate the occurrence of jumps based on the jump detection test of Lee and Mykland (2008) using a confidence level of $\alpha = 5\%$.

except for Tesla Motors and Facebook. The obtained sign of the coefficient is also consistent with the notion that variance swaps offer protection against jumps. Jumps increase the realized variance, and if their impact is sufficiently large, payoffs on variance swaps are positive just as are payoffs on the JUMP factor. The fact that both factors are jointly significant suggests that investors clearly distinguish between pure volatility risk and jump risk and that each of these risks commands a separate premium that contributes to variance swap returns. The average contributions (in absolute terms) of the compensation for jump risk and diffusive risk to variance swap returns appear to be of comparable magnitude. Multiplying the respective regression coefficient by the average risk factor return and taking the average across all underlyings results in a mean contribution of -2.95% for the JUMP factor and -2.34% for the VOL factor. However, even though the JUMP and VOL factors should directly proxy for the two risk factors most commonly thought to be priced in options, the alpha is still highly significant and persistently negative for all underlyings except the CAC40, for which variance swap returns are fully explained by the two factors.

In order to further investigate the explanatory power of the JUMP and VOL factors in the presence of traditional factors, I augment the previous regressions with the Fama and French (1993) factors. The results are shown in Table 8. First, even in the presence of the three traditional factors, the coefficients on both the JUMP and VOL factors remain positive and significant at the 1% level for the majority of regressions, which suggests that the two factors contain infor-

mation that are not already captured by the Fama and French (1993) factors.

Second, for the majority of regressions, the absolute magnitude of the SMB coefficient as well as the total number of regressions for which it is significant is smaller than in Table 5 after the inclusion of the two risk-factor-mimicking portfolios, which points to a link between the SMB factor and either jump or volatility risk or both. Individual regressions with the Fama and French (1993) factors and the VOL and JUMP factor separately show that only the JUMP factor leads to this reduction in the number of significant SMB coefficients. In contrast, for regressions with the VOL factor and without the JUMP factor, the number of significant coefficients on the SMB factor does not change. In this context, the presence of priced jump risk in variance swap returns may also help to rationalize the negative and often significant coefficient on the SMB factor in previous regressions with the Fama and French (1993) factors (c.f. Table 5), that is also documented by Carr and Wu (2009). A possible relation between the SMB factor and jump risk can be substantiated by recent findings of Bollerslev et al. (2015) who examine how individual equity prices are affected by continuous and discontinuous market price movements. For this purpose, they calculate separate betas to measure the continuous and discontinuous comovements. Portfolio sorts on the discontinuous betas lead to statistically significant positive return spreads between the High-Low-portfolio but less so for the continuous beta. Moreover, their results show that particularly smaller firms and stocks with lower book-to-market ratio tend to have higher

Table 7: Regression of variance swap returns on returns on jump-risk mimicking portfolio JUMP and volatility-risk mimicking portfolio VOL

Note: Entries show the OLS estimates and t-statistics (in parentheses) from regressions of continuously compounded 22 day variance swap returns on the 22-day cumulative returns on the jump-risk mimicking portfolio JUMP and the VOL factor. t-statistics are based on standard errors that are adjusted for serial correlation according to the method of Newey and West (1987) with a lag length of 22 days. N denotes the number of total observations.

Underlying	alpha		JUMP		VOL		R ²	N
S&P 500	-0.433	(-16.574)	2.589	(15.635)	2.976	(8.739)	62.93%	2881
Dow Jones Industrial	-0.396	(-11.696)	2.322	(11.512)	2.999	(7.508)	57.72%	1904
NASDAQ100	-0.282	(-9.686)	2.008	(9.625)	2.080	(5.914)	45.08%	2637
Alcoa	-0.233	(-4.142)	1.106	(3.718)	1.376	(1.541)	19.01%	703
Altria (Philip Morris)	-0.338	(-6.209)	0.740	(2.275)	2.480	(3.486)	8.42%	1658
Amazon	-0.222	(-5.173)	1.099	(3.978)	1.608	(2.966)	12.33%	2023
American Express	-0.137	(-3.592)	1.905	(9.443)	2.714	(5.134)	39.49%	1616
Amgen	-0.289	(-7.176)	1.421	(6.846)	1.618	(2.796)	20.04%	1933
Analog Devices	-0.162	(-3.072)	1.325	(5.780)	1.203	(1.802)	18.08%	536
Apple	-0.299	(-7.656)	1.080	(3.976)	1.306	(2.994)	15.39%	1962
Bank of America	-0.152	(-2.652)	1.770	(5.213)	1.210	(2.057)	23.60%	1236
Boeing	-0.213	(-6.337)	1.090	(5.506)	1.390	(2.807)	17.23%	1778
Cisco	-0.197	(-4.056)	1.250	(4.704)	2.699	(4.993)	21.55%	1436
Exxon Mobil	-0.253	(-6.377)	1.557	(6.869)	2.946	(5.880)	37.04%	1565
Facebook	-0.340	(-4.501)	0.345	(0.696)	-3.853	(-2.098)	8.67%	486
General Electric	-0.169	(-4.821)	1.840	(11.456)	1.932	(3.739)	35.11%	1446
Home Depot	-0.265	(-7.672)	1.400	(7.201)	2.308	(5.750)	27.27%	1537
IBM	-0.295	(-8.682)	1.447	(7.250)	2.430	(6.567)	24.33%	2298
Johnson & Johnson	-0.330	(-7.060)	2.101	(6.843)	2.996	(4.689)	34.08%	1058
McDonald's	-0.384	(-9.678)	0.839	(3.929)	2.514	(5.009)	22.97%	1286
Merck	-0.237	(-5.920)	2.012	(8.218)	1.679	(3.194)	23.77%	1663
Metlife	-0.126	(-2.836)	1.669	(4.114)	3.286	(6.193)	47.67%	900
Microsoft	-0.208	(-5.128)	1.213	(5.149)	1.200	(2.482)	14.99%	1863
Monsanto	-0.244	(-6.749)	1.319	(6.487)	2.094	(4.115)	30.19%	1217
Nike	-0.249	(-5.134)	1.165	(4.071)	2.882	(5.091)	21.57%	1241
Pfizer	-0.223	(-5.263)	2.009	(8.171)	2.010	(4.329)	26.51%	1139
Procter & Gamble	-0.357	(-6.448)	1.806	(5.367)	2.757	(6.146)	29.64%	1393
Starbucks	-0.229	(-4.446)	0.729	(3.455)	1.565	(2.006)	9.45%	1316
Tesla	-0.277	(-3.817)	0.096	(0.214)	0.758	(0.781)	0.74%	701
Valero	-0.151	(-3.473)	1.086	(3.482)	0.905	(1.534)	14.11%	1607
Verizon	-0.257	(-6.726)	1.075	(4.184)	2.136	(3.336)	22.76%	1326
WallMart	-0.317	(-7.693)	1.006	(5.173)	2.372	(5.218)	19.25%	1584
AEX	-0.382	(-8.334)	2.496	(4.356)	4.337	(5.929)	61.52%	523
CAC40	0.013	(1.141)	6.140	(30.556)	0.438	(2.296)	90.14%	22
DAX	-0.267	(-6.701)	3.337	(9.152)	3.544	(6.155)	66.22%	940
Euro Stoxx 50	-0.295	(-6.519)	3.266	(10.115)	2.105	(4.067)	64.26%	485
SMI	-	-	-	-	-	-	-	-

discontinuous betas.

This implies that small stocks have a more pronounced tendency to co-jump with the market than larger stocks. [Ang and Chen \(2002\)](#) also document that correlation asymmetries are stronger for smaller stocks and conclude that such stocks are more exposed to common downward movements with the market. Thus, the SMB factor should have a tendency to perform well in good times but less so in bad times when negative jumps occur and variances usually increase. This relation is also supported by the negative correlation of -0.2468 (p-Value 0.0003) between the returns on the JUMP factor and on the SMB factor for the US. As a consequence variance swap returns appear to be negatively correlated with

returns on the SMB factor due to their differential exposure to jumps.

Third, even though the coefficient on the excess market return is still negative and significant at the 1% level for almost all regressions, its magnitude is substantially smaller by an amount of as much as approximately one half compared to the previous regressions without the two additional factors. This finding suggests that information contained in the market excess return is also captured by the JUMP and VOL factors. The negative correlations of -0.2643 (p-Value 0.0003) between returns on the US market portfolio and on the US JUMP factor and of -0.3091 (p-Value 0.0000) between the US market excess return and returns on the VOL fac-

tor for the US further support this view. For one thing, the negative correlations are not surprising since the JUMP factor usually earns positive returns at times when substantial (mostly negative) market moves occur and such moves are often associated with elevated levels of volatility, which in turn affect returns on the VOL factor. However, the underlying economic rationale to which this association can probably be attributed is the notion that a substantial portion of the equity risk premium itself represents compensation for jump risk and volatility risk (e.g. [Bollerslev and Todorov \(2011\)](#)). The intuitive explanation for this relation is that an increased volatility adversely affects investors' investment opportunity sets and that jumps may be associated with tail risk.¹⁴ If investors then feel that the likelihood of a tail event or return volatility increases, they may require a higher return, which leads to an immediate, often substantial, reduction in asset prices. This mechanism is quite similar to the volatility feedback effect explained before. A possible explanation for why elevated volatility levels or perceived fear of extreme events can have important impacts on asset prices and thus on market excess returns is that such factors are suspected to affect aggregate economic output. This is the so-called uncertainty shock channel. For instance, [Bloom \(2009\)](#) finds evidence that uncertainty shocks have substantial effects on the aggregate economy, leading to significant drops in production, output, and productivity growth since higher uncertainty causes firms to delay investments and hiring. While [Bloom \(2009\)](#) relates uncertainty to volatility, [Gourio \(2012\)](#) and [Kelly and Jiang \(2014\)](#) find that perceived tail risk can have similar effects on aggregate economic activity. The changing magnitude of the coefficient on the market excess return after the introduction of the VOL and JUMP factor can therefore be regarded as a further indication that the equity risk premium is linked to jump or tail risk as well as volatility risk.

Finally, even though the alpha coefficient is slightly less negative when the JUMP and VOL factors are included as additional explanatory variables, their overall influence on its magnitude is negligible and it remains significant in all regressions. While it presumably is comprehensible that commonly used risk factors such as those proposed by [Fama and French \(1993\)](#) cannot satisfactorily explain the returns on variance swaps when overall market variance and jump risk constitute independently priced risk factors, it is astonishing that adding the two constructed factor portfolios does not appear to perceptibly reduce the alpha. Especially with regard to the fact that the coefficients on the JUMP and VOL factors behave as one would expect them to do, so that it can reasonably be assumed that these two factors are indeed able to effectively capture the isolated compensation for jump and volatility risk, this is intriguing. Moreover, even if one considers the possibility that a certain portion of single stock variance swap returns represents compensation for the risk

of common changes in idiosyncratic variance, as for example suggested by [Gourier \(2015\)](#), this is unlikely to be the case for highly diversified stock indices such as the S&P 500 for which the alpha is also significant. Consequently, it appears that the pricing of variance swaps is either inefficient, allowing to earn significant abnormal returns that appear to be unrelated to systematic risk factors or that sources of systematic risk other than jump and volatility risk are priced in variance swap rates.

Robustness considerations

Since the assertion that financial assets are inefficiently priced should be made with caution only, it is imperative to take into account possible misspecifications and consequences of transformations of variance swap returns that may have an impact on the results. As previously outlined, the logarithmic version of the variance risk premium, defined as $LVRP_{t,T} = \log(RV_{t,T}/E_t^Q[RV_{t,T}])$, is introduced in order to account for the substantial kurtosis and positive skewness of raw variance swap returns – particularly driven by excessively high, positive returns on long variance swap positions during the recent financial crisis – and thereby make the distribution more Gaussian. The underlying rationale for this transformation is to make the regression residuals more normal. Even though this transformation is also applied by other studies (e.g. [Carr and Wu \(2009\)](#)) and the use of log returns in the context of stock return applications is widely accepted since deviations between continuously compounded returns and discrete returns are typically negligible for small price differentials, its application to variance swap returns comes with potential drawbacks. While both, skewness and kurtosis can indeed be mitigated through the log transformation, the often substantial divergence between realized variance and the variance swap rate can lead to considerable deviations between raw returns ($RV_{0,T}/E_t^Q[RV_{0,T}] - 1$) and their logarithmic version. As a consequence, the logarithmic transformation tilts the mean of the variance swap return distribution disproportionately into the negative domain, which may contribute to the persistently negative alpha.

In order to investigate the impact of the logarithmic transformation on reported results, several robustness regressions are performed. Since the aim is only to provide qualitative insights, the corresponding results can be found in the Online-Appendix.

Table A3 in the Online-Appendix shows the results from robustness regressions of raw 22-day variance swap returns on monthly realizations of the three [Fama and French \(1993\)](#) factors and the two risk factor portfolios JUMP and VOL over the entire sample horizon from January 1996 to July 2015. For raw variance swap returns, the alpha is consistently less negative and significant at the 5%-level for only 12 underlyings including the indices, compared to 33 significant alphas for the logarithmic specification.

In order to further emphasize the influence of the logarithmic transformation, the sample is divided into two subsamples. The first subsample is restricted to cover only the time period from January 1996 to December 2007 and there-

¹⁴For instance, [Bollerslev and Todorov \(2011\)](#) find that the magnitude of the left jump tail which captures fears of dramatic market declines significantly exceeds that of the right-jump tail which is associated with substantial market appreciations.

fore to exclude the time period of unprecedented uncertainty during and following the financial crisis. Since extremely positive as well as extremely negative variance swap returns predominantly occur in later periods, the effect of the logarithmic transformation can be expected to be comparatively moderate in this subsample. Note that this subsample only comprises the US underlyings since the time series for the European underlyings either start later or it is not possible to obtain at least ten valid observations for all regressors (this is the case for the DAX). The second subsample is restricted to cover the time period from January 2008 until July 2015 during which the qualitative and quantitative deviations between raw returns and their logarithmic version can be expected to be more pronounced. Results for the first and second subsample for the same regression specifications as before are reported in Tables A4 and A5 in the Online-Appendix.

Comparing the results for the entire sample and the first subsample, the results differ and deviations between the two specifications are less pronounced for the latter. For the first subsample, the number of negative and significant alphas for the logarithmic specification reduces substantially from 33 to only 15 underlyings whereas the specification with raw returns produces significantly negative alphas for eight underlyings. For both specifications, however, the alphas remain significant for the indices. As expected, qualitative and quantitative differences are more pronounced for the second subsample since the most extreme returns – positive and negative – occur in this period. Regressions with continuously compounded returns again produce 32 significantly negative alphas compared to 12 for regressions with raw returns.

These robustness regressions clearly highlight the importance of the chosen specification of the variance risk premium as well as of the chosen time period and provide evidence that the two specifications might lead to qualitatively different conclusions for certain underlyings when the objective is to determine whether variance swap returns are completely explained by a given set of risk factors. However, even though the number of significant alphas is obviously lower when raw variance swap returns are used, none of the applied regression models is able to fully explain these returns for all underlyings over the entire time horizon or subsets of it. Thus, even though variance swap returns clearly include compensation for jump and volatility risk, these risks are not the only factors that appear to influence the pricing of options and variance swaps, which points toward alternative or additional explanations such as for example model uncertainty (e.g. Drechsler (2013)) or demand-based explanations (e.g. Gârleanu et al. (2009)).

5.2.4. Model uncertainty and uncertainty about macroeconomic fundamentals

In this final section, variance swap returns are related to model uncertainty and uncertainty about macroeconomic fundamentals. The purpose is explicitly not to explain the variance risk premium through risk factors but rather to shed some light on one of the factors that could contribute to the time-varying differences in investors' perception of variance

risk or jump risk that are presumably associated with the substantial fluctuations of the variance risk premium over time (c.f. Figure 1) and to assess whether a measure of model uncertainty could be helpful in explaining variance swap returns in a richer setting than that used here.

The variable to be considered in this setting is motivated by the work of Drechsler (2013) who considers how the variance risk premium is affected when investors face model uncertainty. Due to the model uncertainty, the investor is not only uncertain about a possible outcome but also about the data generating process and thus evaluates his decisions under the "worst case" scenario, in which he particularly overestimates the magnitude and frequency of jumps. Since the investor knows that realized variance will be high if the worst case scenario manifests, his willingness to pay for variance swaps will also be high when model uncertainty is high. Therefore, a direct implication of the model is that the size of the variance risk premium will fluctuate and reflect investors' degree of model uncertainty. In particular, this relation is under consideration in the following.

In order to measure model uncertainty, I follow Drechsler (2013) and use the cross-sectional dispersion (DispGDP), i.e. the variance of implied quarterly growth rates of the nominal gross domestic product (GDP). Consequently, higher values are associated with greater model uncertainty.

The necessary data is taken from the Survey of Professional Forecasters to which the Federal Reserve Bank of Philadelphia provides access.¹⁵ The Survey of Professional Forecasters contains forecasts by individual forecasters for the quarterly level of nominal GDP in billions \$, based on a seasonally adjusted annual rate.

The forecasted growth rate is based on the forecasted level of nominal GDP for the current quarter and the realized level of nominal GDP in the previous quarter. In order to mitigate the effect of outliers, dispersion figures which are more than two standard deviations from the sample mean are removed. Drechsler (2013) argues that information asymmetries can be expected to play only a minor role in the context of forecasting macroeconomic quantities such as GDP because essential information are publicly available. As a consequence, the dispersion of GDP forecasts can naturally be interpreted as the result of different economic models.

Since the individual GDP forecasts are made each quarter based on an adjusted annual growth rate whereas the variance swap returns refer to a period of only 22 trading days, I assume that the DispGDP measure does not only reflect model uncertainty for a particular day but reflects uncertainty around the time the forecasts are made, i.e. at the beginning of a quarter. As a consequence the DispGDP measure at the beginning of a given quarter is matched with the variance swap return over the 22 trading day period starting on the day on which the DispGDP measure is computed. If no such return could be computed on this specific day, I

¹⁵The internet address is: <https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/>

use the next available return calculated on a later day but not later than three calendar days after the day the DispGDP measure is computed. The log of the realized variance over the 22 trading day period ending immediately before the period starts over which the variance swap returns are computed is included as an additional control variable because other studies (e.g. [Bollerslev et al. \(2011\)](#)) document a significantly positive relationship between previously realized volatility and risk-aversion as measured by the variance risk premium. The reason for taking the log of realized variance is that variances have right-skewed distributions whereas logarithmic variances tend to have near Gaussian distributions, which may improve predictions from linear models ([Bekaert and Hoerova \(2014\)](#)).

Since both variables, DispGDP and the preceding variance are measures of uncertainty, including them simultaneously in the regressions could lead to flawed inferences if both measures are highly correlated and convey widely identical information. However, for 23 of the 32 underlyings, the absolute value of the correlation between DispGDP and the previously realized variance is below 0.15 with a maximum correlation of 0.47 and an average value across all underlyings of 0.0306. Thus, the two measures appear to convey information about relatively unrelated sorts of uncertainty.

Table 9 shows the results from regressions of the log variance risk premium on the log of the realized variance during the preceding 22 trading days and DispGDP.

The coefficient on the DispGDP measure is negative for all but two underlyings and significant for some of them albeit not for all suggesting that model uncertainty is indeed associated with a more negative realized variance risk premium. The sign of the coefficient is therefore consistent with the notion that investors are willing to pay more for an instrument that offers insurance against undesirable states when the degree of model uncertainty, i.e. uncertainty about the distribution of possible outcomes, is higher. Even though the regression specification applied here is of course not directly comparable with the model specification of [Drechsler \(2013\)](#), the results are in line with his argumentation that model uncertainty and the associated uncertainty or disagreement about macroeconomic fundamentals may help to explain the magnitude and often substantial fluctuations in the observed variance risk premium. When turning to the realized variance over the preceding 22 trading days, the mostly positive sign on the coefficient contradicts the expected relation that higher variance could be associated with higher risk-aversion in the aftermath. However, the coefficients are mostly insignificant.

The chosen regression specification in this section is relatively parsimonious. As a consequence the results have limited meaning only. However, especially in consideration of the fact that the two factor portfolios that capture jump and volatility risk are not able to fully explain variance swap returns for several underlyings suggests that these returns may reflect compensation for further systematic risks. With regard to this, the findings in this section suggest that model uncertainty may be a potential candidate for such a risk factor whose influence on the pricing of options, and particularly

of variance swaps, could be assessed with more advanced econometric techniques in the future.

6. Conclusion

In this work, I apply a relatively model-free method to extract the expected variance under the risk-neutral measure from prices of traded options and to quantify the realized variance risk premium.

Consistent with prior studies, I find a persistently negative variance risk premium for all stock indices. Moreover, the average return on a long single stock variance swap investment is also negative for all considered underlyings even though the absolute magnitude is often smaller than for the indices and returns exhibit a substantially greater variation. The finding of a negative variance risk premium for single stocks is in line with the results of some studies but contradicts those of others. However, these differential findings may, at least in part, be due to the application of different sample selection criteria since any additional factors and circumstances that systematically affect the cross-section of option prices, such as for example illiquidity discounts, will necessarily be reflected in the estimates of the risk-neutral expected variance and the variance risk premium.

All in all, the persistently negative premium suggests that investors regard variance swaps or the associated option positions as valid instruments to insure against undesirable states and are therefore willing to accept negative average returns on them. Even though the results suggest that continuously selling variance swaps would probably have been profitable for most underlyings over the entire sample period, an analysis of variance swap returns during the last quarter of 2008 reveals that the short side is exposed to considerable risks and the losses can be of excessive magnitude. The maximum returns on synthetic variance swaps during this period amount to several hundred percent for many underlyings.

Apart from quantifying the variance risk premium, the objective of this work is to assess whether return variance constitutes an independently priced risk factor for which the variance risk premium offers compensation.

The analysis reveals that the variance risk premium on the US indices is clearly linked to the return variance on the market portfolio and is more negative when the return variance of the index exhibits stronger covariation with that of the market portfolio. This is indicative of priced market variance risk. In contrast, a similar relation is not discernible for single stocks. However, excess returns on single stock variance swaps can often be fully explained by a systematic variance risk factor proxied by returns on variance swaps with the S&P 500 as underlying which suggests that returns on index as well as on single stock variance swaps are, to a considerable extent, driven by the same factors.

Further analyses show that the classical capital asset pricing model as well as the [Fama and French \(1993\)](#) three-factor model are not able to explain variance swap returns and leave a significantly negative alpha, which further supports the notion of an independently priced market variance risk factor.

Table 9: Regressions of variance swap returns on a measure of model uncertainty

Note: Entries report the OLS-estimates and t-statistics (in parentheses) of non-overlapping regressions of 22-day continuously compounded variance swap returns on a DispGDP and the realized variance over the preceding 22 trading days (pastVariance). t-statistics are not adjusted for serial correlation. N denotes the total number of observations.

Underlying	const		DispGDP		pastVariance		R ²	N
S&P500 Index	-0.391	(-4.599)	-0.385	(-3.927)	2.765	(1.316)	15.33%	70
Dow Jones Industrial	-0.358	(-3.547)	-0.527	(-2.731)	2.876	(0.815)	15.77%	42
NASDAQ100	-0.210	(-2.680)	-0.151	(-1.935)	1.057	(2.189)	8.19%	68
Alcoa	-0.742	(-3.272)	0.409	(1.301)	2.234	(1.538)	22.46%	19
Altria (Philip Morris)	-0.531	(-3.450)	-0.083	(-0.576)	1.632	(1.279)	3.77%	43
Amazon	0.166	(1.395)	-0.064	(-0.629)	-0.024	(-0.445)	0.46%	51
American Express	-0.083	(-0.915)	-0.305	(-2.591)	0.494	(3.560)	13.13%	46
Amgen	-0.099	(-0.804)	-0.484	(-2.726)	0.270	(0.605)	7.25%	46
Analog Devices	0.177	(0.724)	-0.490	(-1.556)	-0.011	(-0.189)	8.89%	15
Apple	-0.231	(-3.075)	0.001	(0.005)	-0.010	(-4.200)	1.50%	51
Bank of America	-0.169	(-1.671)	-0.004	(-0.053)	0.501	(7.752)	20.97%	30
Boeing	-0.159	(-1.989)	0.061	(0.701)	0.052	(2.642)	1.42%	50
Cisco	-0.179	(-1.995)	-0.048	(-0.386)	0.039	(0.703)	1.45%	37
Exxon Mobil	-0.410	(-2.636)	-0.231	(-1.377)	3.062	(0.854)	11.67%	39
Facebook	-0.728	(-5.322)	-0.350	(-1.708)	2.347	(6.695)	57.60%	7
General Electric	-0.001	(-0.011)	-0.218	(-2.283)	-0.560	(-1.890)	11.75%	37
Home Depot	-0.303	(-1.822)	-0.387	(-1.363)	0.966	(1.424)	7.05%	31
IBM	-0.209	(-2.208)	-0.202	(-2.051)	0.672	(0.912)	3.57%	68
Johnson & Johnson	-0.177	(-1.156)	-0.418	(-1.791)	1.203	(1.446)	7.52%	27
McDonald's	-0.340	(-2.793)	-0.728	(-5.236)	1.711	(0.498)	22.15%	26
Merck	0.006	(0.053)	-0.360	(-3.467)	0.561	(1.416)	11.56%	46
Metlife	-0.207	(-1.532)	-0.306	(-1.296)	0.384	(0.536)	4.80%	18
Microsoft	-0.044	(-0.579)	-0.208	(-1.652)	0.032	(1.077)	5.56%	52
Monsanto	-0.460	(-4.929)	-0.260	(-1.221)	1.229	(1.368)	16.64%	30
Nike	-0.267	(-2.039)	-0.041	(-0.138)	-0.048	(-1.208)	0.88%	31
Pfizer	0.064	(0.506)	-0.121	(-0.889)	-0.058	(-7.275)	9.54%	26
Procter & Gamble	-0.109	(-0.897)	-0.673	(-2.801)	0.113	(5.212)	13.81%	38
Starbucks	-0.008	(-0.059)	-0.184	(-1.041)	0.147	(0.261)	2.54%	32
Tesla	-0.547	(-1.277)	-0.383	(-1.248)	0.234	(0.127)	8.12%	12
Valero	-0.073	(-0.671)	0.115	(0.554)	-0.001	(-0.023)	0.60%	36
Verizon	-0.083	(-0.827)	-0.166	(-0.632)	-0.258	(-0.131)	1.41%	31
WalMart	-0.384	(-2.401)	-0.243	(-1.090)	2.793	(3.746)	15.27%	35

In order to gain isolated exposure to the risks associated with the two sources of realized variance, namely continuous and discontinuous price movements, I construct two risk-factor-mimicking portfolios from traded options that are intended to offer exposure to one of these risks while being relatively unaffected by the other. Using these two risk factor portfolios, I find evidence that compensation for jump risk in addition to compensation for diffusive risk represents a substantial part of variance swap returns. The contribution of the two types of risk to average variance swap returns appears to be of comparable magnitude.

Even when the [Fama and French \(1993\)](#) factors are included as additional control variables, the coefficients on the two risk-factor-mimicking portfolios remain positive and highly significant, which further supports the view that jump risk and diffusive volatility risk constitute independently priced risk factors that are not already accounted for by commonly used risk factors.

However, none of the applied regression specifications is able to fully explain observed variance swap returns. Regres-

sion alphas are negative and significant for virtually all underlyings. This circumstance can either result from inefficient pricing of options and variance swaps that allows to earn substantial abnormal returns which are unrelated to systematic risks, or from exposure to additional risk factors that are not considered in the regressions but priced by the market.

Motivated by the work of [Drechsler \(2013\)](#) who links the variance risk premium to model uncertainty and ambiguity aversion, I consider model uncertainty as one such potential factor and find that higher model uncertainty is associated with a more negative variance risk premium. Even though the chosen specification is relatively parsimonious and potentially neglects several important control variables, these results could therefore point toward a variable that may help explain variance swap returns. An interesting extension for future research could therefore be to find a reliable measure of model uncertainty that is available at a relatively high frequency, for example daily or weekly, and test whether this measure helps to explain variance swap returns.

References

References

- Ait-Sahalia, Y. and Lo, A. W. Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2): 499–547, 1998.
- Ait-Sahalia, Y., Cacho-Diaz, J., and Laeven, R. Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics*, 117(3):585–606, 2015a.
- Ait-Sahalia, Y., Karaman, M., and Mancini, L. The term structure of variance swaps and risk premia. *Working Paper*, 2015b.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., and Ebens, H. The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1): 43–76, 2001.
- Anderson, E. W., Ghysels, E., and Juergens, J. L. The impact of risk and uncertainty on expected returns. *Journal of Financial Economics*, 94(2): 233–263, 2009.
- Ang, A. and Chen, J. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494, 2002.
- Bajoux-Besnainou, I. and Rochet, J.-C. Dynamic spanning: Are options an appropriate instrument? *Mathematical Finance*, 6(1):1–16, 1996.
- Bakshi, G. and Kapadia, N. Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, 16(2):527–566, 2003.
- Bates, D. S. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1):181–238, 2000.
- Bates, D. S. The market for crash risk. *Journal of Economic Dynamics and Control*, 32(7):2291–2321, 2008.
- Bekaert, G. and Hoerova, M. The VIX, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2):181–192, 2014.
- Black, F. Studies of stock price volatility changes. In *Proceedings of the 1976 meetings of the Business and Economic Statistics Section, American Statistical Association*, pages 177–181, 1976.
- Black, F. and Scholes, M. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81:637–654, 1973.
- Bloom, N. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.
- Bollerslev, T. and Todorov, V. Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211, 2011.
- Bollerslev, T., Gibson, M., and Zhou, H. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Working Paper*, 2008.
- Bollerslev, T., Gibson, M., and Zhou, H. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics*, 160(1):235–245, 2011.
- Bollerslev, T., Li, S. Z., and Todorov, V. Roughing up beta: Continuous vs. discontinuous betas, and the cross-section of expected stock returns. *Working Paper*, 2015.
- Bondarenko, O. Market price of variance risk and performance of hedge funds. In *Working Paper, University of Illinois at Chicago*, 2004.
- Breedon, D. T. and Litzenberger, R. H. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51:621–651, 1978.
- Britten-Jones, M. and Neuberger, A. Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866, 2000.
- Broadie, M., Chernov, M., and Johannes, M. Understanding index option returns. *Review of Financial Studies*, 22(11):4493–4529, 2009.
- Cao, H. H. and Ou-Yang, H. Differences of opinion of public information and speculative trading in stocks and options. *Review of Financial Studies*, 22(1):299–335, 2009.
- Cao, J. and Han, B. Cross section of option returns and idiosyncratic stock volatility. *Journal of Financial Economics*, 108(1):231–249, 2013.
- Carr, P. and Lee, R. Volatility derivatives. *Annual Review Financial Economics*, 1(1):319–339, 2009.
- Carr, P. and Madan, D. Towards a theory of volatility trading. *Risk Books on Volatility*, (29):417–427, 1998.
- Carr, P. and Wu, L. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.
- Carr, P. and Wu, L. Leverage effect, volatility feedback, and self-exciting market disruptions. *Bloomberg Portfolio Research Paper*, (2009-03), 2011.
- CBOE. The CBOE volatility index - VIX. *Chicago Board Options Exchange, White Paper*, pages 1–23, 2015.
- Coval, J. D. and Shumway, T. Expected option returns. *The Journal of Finance*, 56(3):983–1009, 2001.
- Cox, J. C. The constant elasticity of variance option pricing model. *The Journal of Portfolio Management*, 23(5):15–17, 1996.
- Cremers, M., Halling, M., and Weinbaum, D. Aggregate jump and volatility risk in the cross-section of stock returns. *The Journal of Finance*, 70(2): 577–614, 2015.
- Demeterfi, K., Derman, E., Kamal, M., and Zou, J. A guide to volatility and variance swaps. *The Journal of Derivatives*, 6(4):9–32, 1999.
- Drechsler, I. Uncertainty, time-varying fear, and asset prices. *The Journal of Finance*, 68(5):1843–1889, 2013.
- Drechsler, I. and Yaron, A. What's vol got to do with it. *Review of Financial Studies*, 24(1):1–45, 2011.
- Driessen, J., Maenhout, P. J., and Vilkov, G. The price of correlation risk: Evidence from equity options. *The Journal of Finance*, 64(3):1377–1406, 2009.
- Du, J. and Kapadia, N. The tail in the volatility index. *Working Paper, University of Massachusetts*, 2011.
- Dupire, B. Pricing with a smile. *Risk*, 7(1):18–20, 1994.
- Ellsberg, D. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75:643–669, 1961.
- Engle, R. and Figlewski, S. Modeling the dynamics of correlations among implied volatilities. *Review of Finance*, 2014.
- Fama, E. F. and French, K. R. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.
- Fama, E. F. and French, K. R. The capital asset pricing model: Theory and evidence. *The Journal of Economic Perspectives*, 18:25–46, 2004.
- Fournier, M. and Jacobs, K. Inventory risk, market-maker wealth, and the variance risk premium: Theory and evidence. *Rotman School of Management Working Paper*, (2334842), 2015.
- Gârleanu, N., Pedersen, L. H., and Potoshman, A. M. Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299, 2009.
- Glosten, L. R., Jagannathan, R., and Runkle, D. E. On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801, 1993.
- Gourier, E. Pricing of idiosyncratic equity and variance risks. *Working Paper Princeton University*, 2015.
- Gourio, F. Disaster risk and business cycles. *American Economic Review*, 102: 2734–2766, 2012.
- Green, T. C. and Figlewski, S. Market risk and model risk for a financial institution writing options. *The Journal of Finance*, 54(4):1465–1499, 1999.
- Herskovic, B., Kelly, B., and Lustig, H. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *National Bureau of Economic Research Working Paper 20076*, 2014.
- Heston, S. L. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.
- Hull, J. *Options, futures and other derivatives*. Pearson/Prentice Hall, 2009.
- Hull, J. and White, A. The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300, 1987.
- Hull, J. and White, A. OIS discounting, interest rate derivatives, and the modeling of stochastic interest rate spreads. *Journal of Investment Management*, 13:64–83, 2015.
- Ilmanen, A. Do financial markets reward buying or selling insurance and lottery tickets? *Financial Analysts Journal*, 68(5):26–36, 2012.
- Jiang, G. J. and Tian, Y. S. The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342, 2005.
- Jiang, G. J. and Tian, Y. S. Extracting model-free volatility from option prices: An examination of the VIX index. *The Journal of Derivatives*, 14(3):35–60, 2007.
- Kelly, B. and Jiang, H. Tail risk and asset prices. *Review of Financial Studies*, 27(10):2841–2871, 2014.
- Lee, S. S. and Mykland, P. A. Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies*, 21(6):2535–2563, 2008.
- Merton, R. C. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.
- Neuberger, A. The log contract. *The Journal of Portfolio Management*, 20: 74–80, 1994.
- Newey, W. K. and West, K. D. A simple, positive semi-definite, heteroskedas-

- ticity and autocorrelation consistent covariance matrix. *Econometrica*, 55: 703–708, 1987.
- Pan, J. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50, 2002.
- Ross, S. A. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360, 1976.
- Schürhoff, N. and Ziegler, A. Variance risk, financial intermediation, and the cross-section of expected option returns. *CEPR Discussion Paper No. 8268*, 2011.
- Staum, J. Incomplete markets. In Birge, J. R. and Linetsky, V., editors, *Handbooks in operations research and management science*, chapter 12. Elsevier, 2007.
- Todorov, V. Variance risk-premium dynamics: The role of jumps. *Review of Financial Studies*, 23(1):345–383, 2010.
- Wu, G. The determinants of asymmetric volatility. *Review of Financial Studies*, 14(3):837–859, 2001.