Essays on the Interaction of Option and Equity Markets

D I S S E R T A T I O N
of the University of St.Gallen,
School of Management,
Economics, Law, Social Sciences,
and International Affairs
to obtain the title of
Doctor of Philosophy in Finance

submitted by

Alexander Feser

from

Germany

Approved on the application of

Prof. Dr. Manuel Ammann

and

Prof. Paul Söderlind, PhD

Dissertation no. 4946

Difo-Druck GmbH, Untersiemau 2020

The University of St.Gallen, School of Management, Economics, Law, Social Sciences and International Affairs hereby consents to the printing of the present dissertation, without hereby expressing any opinion on the views herein expressed.

St.Gallen, October 22, 2019

The President:

Prof. Dr. Thomas Bieger

To Laura and my parents

Contents

List of Figures

List of Tables

Abstract

How do option and equity markets interact with each other? This is the central question that is answered from three different angles in this dissertation.

The first Chapter discusses how option-implied information is incorporated into equity markets. Based on a novel rescaled option-implied Value-at-Risk (rVaR) measure, it is shown that option-implied information is priced differently depending on whether it is based on options with strikes close to the current price of the underlying or far-out-of-the-money options. The findings provide novel insights in the joint interaction between option and equity markets and help to explain contradictory results in previous studies.

The second chapter provides an in-depth analysis of how to estimate risk-neutral moments robustly. A simulation and an empirical study show that estimating risk-neutral moments presents a trade-off between (1) the bias of estimates caused by a limited strike price domain and (2) the variance of estimates induced by micro-structural noise. The best trade-off is offered by option-implied quantile moments estimated from a volatility surface interpolated with a local-linear kernel regression and extrapolated linearly.

The third chapter expands volatility targeting to option strategies. The chapter shows that option trading strategies can be managed by increasing exposure if volatility is low and reducing exposure if volatility is high to achieve a constant risk exposure over time. These volatility controlled option strategies generate economically and statistically significant alphas over their unmanaged counterparts, have reduced maximum drawdowns, lower downside risk, and more normal return distributions.

Zusammenfassung

Wie interagieren Options und Aktienmärkte miteinander? Die vorliegende Dissertation beantwortet diese Frage aus drei verschiedenen Blickwinkeln.

Das erste Kapitel diskutiert wie Informationen aus Optionsmärkten in Aktienmärkte inkorporiert werden. Basierend auf einem neuen Options-implizierten Value-at-Risk (rVaR) Mass, wird gezeigt, dass Information aus Optionsmärkten anders in Aktienmärkte eingepreist wird abhängig davon ob rVaR von Optionen mit Ausübungspreisen nahe des aktuellen Aktienpreises oder von Optionen die weit aus dem Geld sind geschätzt wird. Die Ergebnisse liefern neue Erkenntnisse über die gemeinsame Interaktion von Options und Aktienmärkten und helfen widersprchliche Ergebnisse aus vorherigen Studien zu erklären.

Das zweite Kapitel enthält eine detaillierte Analyse der Messmethoden von risikoneutralen Momenten. Eine simulations und eine empirische Studie zeigen, dass die Schätzung risiko-neutraler Momente einen Kompromiss darstellt zwischen (1) der Verzerrung von Schätzungen, die durch einen begrenzten Ausübungspreisbereich verursacht wird und (2) der Varianz von Schätzungen, die durch Mikrostrukturrauschen verursacht wird. Der beste Kompromiss wird durch Options-implizierte Quantilmomente geboten, die von einer Volatilitätsoberfläche geschätzt werden, die mit einer lokal-linearen Kernelregression interpoliert und linear extrapoliert wird.

Das dritte Kapitel zeigt, dass Volatilitätsmanagement auch auf Optionsstrategien anwendbar ist. Das Risiko von Optionshandelsstrategien kann kontrolliert werden, indem das Investment in eine Strategie erhöht wird, wenn die Volatilität niedrig ist und das Investment reduziert wird, wenn die Volatilität hoch ist. Diese volatilitätsgemanageten Optionsstrategien generieren ökonomisch und statistisch signifikante Alphas, haben kleinere maximale Drawdownsx und normalere Renditeverteilungen.

Synopsis

Equity and Option markets are inherently connected and in a frictionless perfect market both equity and option markets would contain the same information. However, in reality equity and option markets contain different sets of information and interact with each other. In this dissertation the interaction of equity and option markets is analyzed from three different perspectives.

The first chapter of the dissertation analyzes the interaction between single stock option markets and their underlying equity returns. The chapter aims to unify prior contradictory results on the joint cross-section of option and equity markets. The results show that option-implied rescaled Value-at-Risk is incorporated into equity markets differently depending on whether it is based on options with strikes close to the current price of the underlying or far-out-of-the-money options. If the rVaR is estimated from options close-to-the-money, i.e., the 50% rVaR, stocks with high risk outperform stocks with low risk by 0.60% per month. This finding is in line with downside risk-averse investors and confirms previous findings of, e.g., Conrad, Dittmar, & Ghysels (2013). In contrast, if rVaR is estimated from far-out-of-the-money options, i.e., the 90% rVaR, stocks with high risk underperform stocks with low risk by 0.42% per month, implying that stocks with low risk have higher returns in the cross-section of returns. These findings are consistent with Rehman & Vilkov (2012), Schneider, Wagner, & Zechner (2015), and Stilger, Kostakis, & Poon (2016). Our results show that it is important where option-implied risk measures are measured and are consistent with investors who prefer reliable information over unreliable information and explain contradictory results of prior studies.

The second chapter of the dissertation complements the first chapter by providing an in-depth analysis of how to estimate risk-neutral moments robustly. Information embedded in option prices is valuable for practitioners, regulators and academics alike.

Unsurprisingly, an extensive literature has developed around the use of option-implied information, but there is surprisingly little literature which rigorously examines the efficiency of different risk-neutral moment estimators. The second chapter of the dissertation closes this gap in the literature. In theory, the estimation of risk-neutral moments from a continuum of option prices is a straight forward task. However, the practical estimation of risk-neutral densities from empirical option data is subject to many biases caused by discrete option strike prices and micro-structural noise. The second chapter performs a horse race among different estimation techniques for option-implied central and quantile moments. The results show that risk-neutral moment estimation is subject to a trade-off between the bias of estimates caused by a limited strike price domain and the variance of estimates induced by micro-structural noise. However, two tested methods offer especially favorable bias-variance trade-offs. First, model-free quantile moments estimated from a local-linear kernel regression with linear extrapolation have only a small bias even under narrow domains and are only moderately affected by micro-structural noise. Second, central risk-neutral moments based on a cubic smoothing spline with horizontal extrapolation are not strongly affected by micro-structural noise and their bias under narrow domains is acceptable.

In the third chapter the interaction between index option markets and their underlying index is analyzed. In contrast, the first chapter analyzed the interaction for single stock options. Investment strategies that combine option and equities, such as covered call or put-write strategies, are popular investments. They offer returns comparable to an investment into the underlying at a lower volatility thus generating high Sharpe ratios (Israelov & Nielsen, 2015a). However, put- and buy-write strategies also carry high levels of downside risk which materializes in occasional crashes. The third chapter shows that option trading strategies can be managed by increasing exposure if volatility is low and reducing exposure if volatility is high to achieve a constant risk exposure over time. These volatility controlled option strategies generate

economically and statistically significant alphas over their unmanaged counterparts, have reduced maximum drawdowns, lower downside risk, and more normal return distributions. The chapter also analyzes the drivers of the success of volatility control in option strategies. Option trading strategies expose investors to the equity risk premium and the variance risk premium (Coval & Shumway, 2001). Using the decomposition of Israelov & Nielsen (2015a), the chapter shows that the gains from volatility control are the most statistically significant for the variance risk premium return.

1 Option-Implied Value-at-Risk and the Cross-Section of Stock Returns

Manuel Ammann & Alexander Feser

Status: Published in the Review of Derivatives Studies

ABSTRACT

Based on a novel rescaled option-implied Value-at-Risk (rVaR) measure, we show that option-implied information is priced differently depending on whether it is based on options with strikes close to the current price of the underlying or far-out-of-the-money options. If the rVaR is estimated from options close-to-the-money, i.e., the 50% rVaR, stocks with high risk outperform stocks with low risk by 0.60% per month, in line with downside risk-averse investors. In contrast, if rVaR is estimated from far-out-of-the-money options, i.e., the 90% rVaR, stocks with high risk underperform stocks with low risk by 0.42% per month, implying that stocks with low risk have higher returns in the cross-section of returns. Our results are consistent with investors who prefer reliable information over unreliable information and explain contradictory results of prior studies.

Keywords: option-implied moments, option-implied skewness, downside risk

JEL classification: C14, G11, G12, G13, G14.

We thank Yakov Amihud, Fred Benth, Peter Carr, Magnus Dahlquist, Ralf Elsas, Joachim Grammig, Fabian Hollstein, Kjell Nyborg, Jens Jackwerth, Joël Peress, Marcel Prokopczuk, Stefan Ruenzi, Nic Schaub, Christian Schlag, David Schimko, Paul Söderlind, Erik Theissen, Grigory Vilkov, as well as the participants of the NYU Tandron Finance seminar, the annual conference of the Swiss Society for Financial Market Research 2017, of the 23rd annual meeting of the German Finance Association 2016, of the doctoral workshop at the 23rd annual meeting of the German Finance Association 2016, the Brown Bag Seminar of the University of St.Gallen, the 2016 joint seminar session of the University of St.Gallen and the University of Konstanz, and the 2015 Topics in Finance Seminar in Davos for helpful discussions and comments. All errors are our own.

1.1 Introduction

A vast literature documents the predictive power of option-implied information for equity markets (e.g. Xing, Zhang, & Zhao, 2010; Rehman & Vilkov, 2012; Conrad et al., 2013; An, Ang, Bali, & Cakici, 2014; Stilger et al., 2016; Schneider et al., 2015). Consequently, option-implied information is commonly used by practitioners and academics alike. However, relatively little is known about the source of this return predictability.

We document that the pricing of option-implied information is dependent on the moneyness of the underlying options. In particular, we find that option-implied information based on options with strikes close to the current price of the underlying generates economically sensible premia in the cross-section of returns, i.e., stocks with higher implied risk have higher returns than stocks with low risk. In contrast, we find a market anomaly if option-implied information is obtained from far-out-of-the-money options. In that case, stocks with low implied risk have higher returns than stocks with high risk.

Our contribution to the literature is twofold: First, we introduce option-implied rescaled Value-at-Risk (rVaR). Value-at-Risk (VaR) is a popular risk-measure among practitioners and academics as it is easy to interpret and compute. VaR is essentially a quantile in the left tail of the distribution of stock returns and can thus also be obtained from the risk-neutral distribution. The reliance on quantiles to compute VaR makes it inherently robust to noisy and sparse data, two challenges frequently encountered when working with option data. More importantly, the VaR at a certain confidence level depends only on the options relatively close to the corresponding quantile. For example, a 30 day 90% VaR is estimated from options far in the left tail of the risk-neutral distribution. In contrast, a 30-day 50% VaR is estimated from options close to the current strike price. This offers two advantages: First, rVaR measures are local tail risk measures allowing us to measure tail risk from a spe-

cific set of options. Second, their local construction ensures that rVaR measures are not affected by the truncation errors that impact VIX and option-implied skewness, which are based on infinite sums of options Barone Adesi (2016). Moreover, another major advantage of risk-neutral VaR measures over risk-neutral central moments is that they are less affected by the change of measure at short horizons making the risk-neutral VaR is closer to the real-world VaR. Barone Adesi (2016). We rescale VaR because it is strongly dependent on the level of volatility. Rescaling is necessary as option-implied volatility has seasonal patterns, for example before the announcement of earnings and major news (e.g. Jin, Livnat, & Zhang, 2012; Lucca & Moench, 2015).

Second, and more importantly, we contribute to the literature by showing that option-implied information has two distinct pricing effects depending on the moneyness of the underlying options. To the best of our knowledge, this is the first study that documents that information from the option markets is not incorporated homogeneously into the equity markets. Instead, investors have informational preferences for information obtained from options with reliable price data close-to-the-money while disregarding information from far-out-of-the-money options. In our study, we analyze the pricing of option-implied information based on rVaR. If the rVaR is estimated from options close-to-the-money, stocks with higher implied risk outperform stocks with lower implied risk. For example, a portfolio that is long in the upper tercile of stocks sorted on the 50% rVaR and short in stocks in the lower tercile yields a monthly return (t-statistic) of 0.60% (4.51) per month. This finding is in line with downside-risk averse investors (e.g. Lettau, Maggiori, & Weber, 2014). In contrast, if rVaR is estimated from far-out-of-the-money options, the pricing effect reverses and stocks with low implied risk outperform stocks with high implied risk. For example, a High-minus-Low (HML) portfolio sorted on the 90% rVaR yields a monthly return (t-statistic) of -0.42% (2.90) per month.

Our findings are consistent with investors who prefer information from reliable

sources over information embedded in unreliable sources, such as options far-out-of-the-money. Options with strike prices close to the current price of the underlying are often relatively liquid and their prices are therefore considered reliable. In contrast, far-out-of-the-money options are often illiquid and their prices can be stale. Therefore, information obtained from far-out-of-the-money options can be viewed as less reliable and is potentially neglected by investors, despite containing valuable information. Research by Chakravarty, Gulen, & Mayhew (2004) and Augustin, Brenner, & Subrahmanyam (2015) shows that far-out-of-the-money options contain valuable information. Chakravarty et al. (2004) and Augustin et al. (2015) argue that informed traders prefer to trade in out-of-the-money options as they offer greater leverage. Some information is thus incorporated into far-out-of-the-money option prices before it is incorporated into the stock market. This explanation is consistent with the findings of Stilger et al. (2016) and Gkionis, Kostakis, Skiadopoulos, & Stilger (2018). Stilger et al. (2016) show that the anomaly is particularly pronounced among overvalued stocks that are difficult to sell short. If an insider possesses information about the true (lower) valuation of such a company, buying far-out-of-the-money put options is a viable strategy as they offer greater leverage and protect the investor from margin calls before the true valuation is released, e.g. at the next earnings announcement. Our argumentation hinges on the idea that investors neglect information in far-out-of-the-money options even though valuable information is contained in their prices. The observed market anomaly of stocks sorted on the 90% rVaR is consistent with a combination of exploitable information in far-out-of-the-money options and the informational preference for close-to-the-money information.

Our hypothesis that the findings are driven by investors that prefer reliable over unreliable information implies that the pricing effect should be largest if the information from close-to-the-money and far-out-of-the-money options is contradictory, i.e. the 50% rVaR implies a high risk and the 90% rVaR implies a low risk. We test this implication of our hypothesis in a double sort on the 50% and 90% rVaR. The

results show that the diagonal High-Low minus Low-High portfolio yields a monthly return of 1.14% with a t-statistic of 4.83. In contrast, if stocks with a low disagreement between close-to-the-money and far-out-of-the-money options are traded, i.e the High-High minus Low-Low diagonal portfolio, the return decreases to 0.35% per month.

The two different pricing effects of option-implied information are distinct from each other. We find that the pricing of risk measures based on far-out-of-the-money options, such as the 90% rVaR, is persistent for up to 6 months. In contrast, the majority of the return of risk measures based on options close-to-the-money occurs within 5 days of portfolio formation. This finding shows that the information from close-to-the-money options is incorporated into the prices of the underlying at a much faster rate.

Furthermore, our results provide an explanation for contradictory results on the pricing of option-implied skewness in previous studies. Investors are downside risk-averse (e.g. Lettau et al., 2014). Therefore, stocks that exhibit a higher level of downside risk than upside potential (left-skewed) should carry a premium in the cross-section of returns. Even though there is plenty of evidence that stocks with more negatively skewed historical returns carry a risk premium (e.g., Kraus & Litzenberger, 1976; Harvey & Siddique, 2000), the literature does not agree on the premium of stocks with left-skewed risk-neutral distributions. Conrad et al. (2013) find that left-skewed stocks have a higher return than right-skewed stocks. However, Rehman & Vilkov (2012), and Stilger et al. (2016) find that left-skewed stocks have significantly lower returns than right-skewed stocks. These contradictory findings constitute a puzzle of risk-neutral skewness pricing. All studies rely on the option-implied moments of Bakshi, Kapadia, & Madan (2003) but implement them differently. The study of Conrad et al. (2013) relies to a larger extend on traded and relatively liquid options. In contrast, the studies of Rehman & Vilkov (2012) and Stilger et al. (2016)

extrapolate the implied volatility surface and thus rely to a larger extent on the information content of far-out-of-the-money options. The findings of these studies are consistent with our results. Conrad et al. (2013) appear to measure skewness from close-to-the-money options and thus find an effect consistent with the pricing of 50% rVaR. In contrast, Rehman & Vilkov (2012) and Stilger et al. (2016) find an effect which is consistent with the pricing of 90% rVaR.

Our work is of economic importance because it shows how information from option markets is incorporated into equity prices. If investors receive information from multiple options, they will focus on information obtained from more reliable options. Therefore, even if information from far-out-of-the-money options exists, investors will still prefer to rely on information from options closer to the money. This behavior allows market inefficiencies, such as in Rehman & Vilkov (2012), Schneider et al. (2015), and Stilger et al. (2016), to persist even in a sample of S&P 500 stocks. This preference for reliable information is therefore a soft limit to arbitrage persisting even among large stocks. Traders searching for arbitrage opportunities should therefore attempt to trade on competing information where one source of information is relatively more reliable, but both sources contain valuable information. Previous work in this area delivers explanations for the existence of the momentum or the book-to-market effect. However, in contrast to our study, past studies usually assume that information uncertainty is related to firms. For example, G. Jiang, Lee, & Zhang (2005) and Zhang (2006) provide evidence that the momentum effect is particularly pronounced for firms with high information uncertainty. Daniel & Titman (2006) show that the book-to-market value ratio generates return predictability because it proxies for the past intangible return of a company which is not priced.

Moreover, we contribute to the literature on the pricing of option-implied information by providing an explanation for contradictory results of previous studies. Conrad et al. (2013) sort portfolios on the central risk-neutral moments of Bakshi et al. (2003)

and find that low skewness stocks earn a premium. They motivate their findings by the theoretical model of Barberis & Huang (2008) who argue that investors with cumulative prospect theory preferences demand securities with right-skewed payoffs (lottery stocks). Bali & Murray (2013) create skewness assets from portfolios of options which are statically delta and vega hedged and find that skewness is negatively related to returns. Therefore, their approach is related to the skewness risk premium of Kozhan, Neuberger, & Schneider (2013). Rehman & Vilkov (2012) and Stilger et al. (2016) sort portfolios on the risk-neutral moments of Bakshi et al. (2003) estimated from an extrapolated implied volatility smile but find that high skewness stocks carry a premium. Both studies show that option traders are able to identify overvalued stocks and that short-sale constraints increase the effect of risk-neutral skewness on future stock returns, consistent with a limits of arbitrage argument (Shleifer & Vishny, 1997). Gkionis et al. (2018) extend the results of Stilger et al. (2016) and find that risk-neutral skewness also captures upside information in stocks. Schneider et al. (2015) confirm the results of Rehman & Vilkov (2012) and Stilger et al. (2016). They show theoretically that skewness is linked to credit risk and to low risk anomalies such as betting-against-beta (Frazzini & Pedersen, 2014) or the low volatility anomaly (Ang, Hodrick, Xing, & Zhang, 2006, 2009). Xing et al. (2010) use the slope of the implied volatility surface as a proxy for skewness and find that stocks with the steepest slopes underperform stocks with flat slopes. They motivate their results by informed trading and demand-based option pricing as in Garleanu, Pedersen, & Poteshman (2009).

The paper proceeds as follows. Section 1.2 presents the data used in the study. Section 1.3 introduces rVaR and describes the estimation procedure from option data. Section 1.4 presents our main empirical results. Section 1.5 analyzes the robustness of our results. Section 1.6 concludes.

1.2 Data

Daily option data of all S&P500 constituents is sourced from the OptionMetrics databases (provided by Wharton Research Data Service) for the sample time period from January 1996 to August 2015. The Volatility Surface database contains implied volatilities with fixed maturities and fixed deltas. Implied volatilities are reported for (absolute) deltas of 0.20, 0.25, 0.30, 0.35, 0.40, 0.45, 0.50, 0.55, 0.60, 0.65, 0.70, 0.75 and 0.80. OptionMetrics only reports implied volatilities if there is sufficient underlying option data to compute it. Our results are based on a maturity of 30 days. Volume and open interest data of options is also sourced from the OptionMetrics volume database. Option data is downloaded for the S&P 500 constituents from 1996-2015 on a daily basis.

Price and return data is downloaded from the Center for Research in Security Prices (CRSP). The CRSP value-weighted US equity market index is used as a market proxy. Risk factors of the Fama & French (2016) five factor model are downloaded from Kenneth French's website. S&P500 constituents are obtained from Compustat. We only include a particular stock-month combination in our sample if the stock is part of the S&P500 at the time of portfolio creation.

1.3 Estimating option-implied Value-at-Risk

The extraction of option-implied Value-at-Risk is similar to the extraction of risk-neutral densities and has recently been outlined by Barone Adesi (2016). The VaR of a stock is defined as $q(1-p)$ where q is the quantile function and p is the desired confidence level. To find the VaR at a certain level p, we need to obtain the quantile function which is the inverse of the cumulative distribution function (CDF). The seminal paper of Breeden & Litzenberger (1978) outlines the theory to obtain the CDF. The primary reason for the popularity of the Breeden-Litzenberger theorem is that it does not make an assumption about the price process of stocks and therefore

allows recovery of the CDF in a model-free way as:

$$F(S_T < K) = e^{r\tau} \frac{\delta P}{\delta K} \tag{1}$$

$$\tag{2}$$

where P is the price of a put, r is the risk-free rate, τ is the time to maturity, and F is the CDF of the underlying under the risk-neutral measure. Hence, the CDF of the option-implied price distribution can be obtained by estimating the first order derivative of a put option with respect to the strike price. Once the CDF is obtained it can be inverted.

While the Breeden-Litzenberger theorem provides an elegant theoretical approach to obtain the risk-neutral distribution, it requires a continuum of option prices to compute the first order derivative. However, real-world option data of single stock equities is often sparse, spanning a narrow domain with large gaps between strike prices. Therefore, estimating the risk-neutral cumulative density function reliably requires to interpolate between the known strike prices, and to extrapolate outside of the observed strike-price domain. We interpolate between strikes with a cubic smoothing spline and extrapolate horizontally as in Carr & Wu (2009) or Rehman & Vilkov (2012) which is the standard practice in the literature. Horizontal extrapolation is popular as it is a simple and robust method that leads to relatively stable estimates over time. Moreover, horizontal extrapolation of the volatility smile assumes that the volatility smile is flat beyond the observed strike-price domain. This assumption thus effectively adds the tails of a normal distribution to the observed part of the risk-neutral distribution. To make sure that our results do not depend on this assumption, we also present a different implementation in Section 1.5.

Our estimation strategy is as follows: We begin by filtering out all options with an absolute delta larger than 0.55 and options with negative implied volatility. Based on

the remaining options, we proceed by interpolating the implied volatility surface with a cubic spline and extrapolate the volatility smile horizontally over a log-moneyness range from -5 to $+5$ with a spacing of 1 cent between strikes. It is important to note that options of single equities are American options, but the Breeden & Litzenberger (1978) theorem holds for European options only. We follow the standard approach in the literature and convert the implied volatility smile of American options into European option prices (e.g., Carr & Wu, 2009; Rehman & Vilkov, 2012). Tian (2011) analyzes the bias introduced by this procedure and finds that the bias is small for out-of-the-money options. Moreover, using American option data introduces another rarely discussed challenge. The put-call parity only holds for European options and thus implied volatilities of American put and call options close-to-the-money often diverge. Simply joining the implied volatility at-the-money will introduce a jump in most cases which leads to instable results. To avoid this issue, we overlay the implied volatilities of put and call options by filtering out options with an absolute delta larger than 0.55 and not, as usually, by filtering out in-the-money options. We differentiate the option prices obtained from our inter- and extrapolation scheme numerically by centered differences. This provides us with a dense set of point estimates of the CDF. We convert strikes to the natural logarithm of moneyness to ensure that we are estimating return distributions and not price distributions. Moreover, in some cases the estimated CDF is not monotonic leading to negative PDFs. We follow the suggestion of Aït-Sahalia & Lo (1998) and run an isotonic regression on the CDF to ensure it is monotonic. Ensuring monotonicity is important as horizontal extrapolation sometimes leads to non-monotonic CDFs at the end points of the observable strike prices. Finally, we interpolate our dense set of point estimates linearly to find the exact location of the quantiles from 0.05 to 0.95 in 0.05 steps and estimate option-implied rVaR as:

$$rVaR(p) = -\frac{q(1-p)}{q(0.75) - q(0.25)} \tag{3}$$

where q is the quantile function. Option-implied rVaR is the rescaled Value-at-Risk. It is necessary to rescale the VaR to remove the influence of changes in volatility. An increase in volatility automatically increases the rVaR as the distribution is stretched. The influence of volatility on VaR is particularly large for quantiles that are further in the tail. For example, the correlation between the implied 30-day 5% VaR and Bakshi et al. (2003) (BKM) implied volatility is 0.93. In contrast, the correlation between implied 30-day 5% rVaR and BKM implied volatility is -0.004. To make sure that our results are not driven by patterns in volatility, such as the volatility drift before news announcements (e.g. Jin et al., 2012; Lucca & Moench, 2015), we rescale VaR. Note that we choose to rescale with the inter-quartile range to make our measure more robust to outliers. However, our results are unaffected by the choice of the scaler and in Section 1.5 we also provide results scaled with option-implied volatility. We also multiply the rVaR with -1 and therefore a higher rVaR corresponds to a higher tail risk. The estimation procedure is perfomed daily for all stocks in our sample, but we keep only the last observation in each month. Furthermore, we choose to omit the maturity of rVaR and write e.g. 90% rVaR instead of 30-day 90% rVaR to increase readability. All results are based on options with a maturity of 30 days.

rVaR is a scale free measure of downside risk, i.e. it measures how far a certain quantile is shifted to the left. A high rVaR thus corresponds to a high risk. The interpretation is similar to that of the regular Value-at-Risk. For example, the 30-day 95% rVaR indicates the maximum rescaled loss with 95% certainty over a period of 30 days. However, unlike the classic VaR, rVaR provides a volatility invariant measure of tail risk, such as skewness or kurtosis. After the estimation and merging of data we are left with 926 firms which were members of the S&P 500 during the sample period with an average of 468 firms per month. The firms come from a broad spectrum of industries with 110'467 firm-month combinations. We also estimate the moments of Bakshi et al. (2003) from the same volatility smile.

Table 1: Summary statistics of option-implied rVaR

This table displays descriptive statistics of option-implied rescaled Value-at-Risk (rVaR) estimated from 30 day options at different quantiles. Option-implied rVaR is defined as $-\frac{q(1-p)}{q(0.75)-q(0.25)}$ where q is the quantile function and p is the desired percentile. For example, the 95% rVaR corresponds to the common 30 day 95% Value-at-Risk rescaled with the interquartile range. The sample contains 926 stocks and spans the period between January 1996 and August 2015 with a total of 110'467 month-company combinations.

	Mean	*Std.*	*min*	25%	50%	75%	*max*
50% rVaR	-1.74	2.98	-9.40	-3.71	-1.98	0.22	5.02
55% rVaR	0.79	3.23	-7.39	-1.66	1.10	3.15	8.00
60% rVaR	3.71	3.77	-5.23	1.46	4.37	6.31	11.89
65% rVaR	7.33	4.29	-3.55	5.35	8.10	10.11	15.93
70% rVaR	11.43	4.54	-2.63	9.52	12.14	14.38	20.11
75% rVaR	15.79	4.85	-1.00	13.72	16.53	18.87	24.94
80% rVaR	21.76	6.02	2.69	18.47	22.65	26.59	30.90
85% rVaR	30.74	6.37	5.83	30.55	32.64	34.17	37.81
90% rVaR	39.92	5.26	10.16	38.79	40.70	42.42	47.56
95% rVaR	52.15	3.70	41.13	49.95	52.14	54.26	63.67

Summary statistics for option-implied moments are provided in Table 1. The mean of the rVaR increases from -1.74% to 52.15% as the confidence level increases from 50% to 95%. Our rVaR measure contains no visible outliers and the standard deviation of all confidence levels is low compared to the mean. For example, the standard deviation of the 95% rVaR is 3.70% with a mean of 52.15%.

1.4 Empirical results

In this section we present our main empirical findings. We document that the pricing of option-implied information is dependent on the moneyness of the options leading to two distinct pricing phenomena. Our results are consistent with investors that prefer reliable over unreliable information. Consequently, investors rely on information obtained from options close-to-the-money and disregard information from

far-out-of-the-money options. Our findings are also consistent with the observed puzzle of risk-neutral skewness and deliver an explanation for the contradictory results of previous studies (e.g., Conrad et al., 2013; Stilger et al., 2016).

1.4.1 Option-implied rVaR

At the end of the last trading day of each month all current members of the S&P 500 are sorted into three tercile portfolios based on their 30 day option-implied rVaR with varying confidence levels. We begin with the 95% rVaR and decrease the confidence level in 5% steps until we reach the 50% rVaR. The portfolio sorts are independent of each other and each sort is performed on the entire sample. With this strategy we are able to capture information content from different areas of the implied volatility surface because rVaR is only dependend on options close to the particular percentile. A higher confidence level of rVaR is estimated from smaller percentiles, i.e. the 95% rVaR is based on the 5th percentile, while the 50% rVaR is based on the 50th percentile. Consequently, higher confidence levels are estimated from put options that are further out-of-the-money. We use breakpoints of 0.33 and 0.67.

The results in Table 2 show the portfolio returns over the next month. The average return of the equal-weighted HML portfolio increases from -0.37% per month to 0.61% per month as the confidence level decreases from 95% to 50%. Furthermore, the t-statistic changes its sign from -2.37 to 3.95. The results uncover two clearly distinct pricing effects. Stocks sorted on rVaRs with confidence levels larger than 80% have a significantly negative HML return, i.e. stocks with lower risk carry a higher return which is inconsistent with downside risk-averse investors. In contrast, stocks with confidence levels smaller than 75% have a significantly positive HML portfolio return, i.e. stocks with higher risks carry a higher return which is consistent with downside risk-aversion. This finding is interesting because it shows that information from the option market is not incorporated homogeneously into the equity markets.

Instead, information from the center of the risk-neutral distribution is incorporated differently than information from the tail of the risk-neutral distribution. Prices of options close-to-the-money are usually considered reliable as they have low bid-ask spreads and are liquid. In contrast, information obtained from non-traded options priced by extrapolating the implied volatility smile is considered as unreliable and is therefore not incorporated into option prices. The returns of the HML portfolio are sizable, a 90% rVaR sorted HML portfolio generates a monthly return of -0.36%, while an 50% rVaR sorted portfolio generates a monthly return of 0.61%.

The observed effect becomes even stronger for value-weighted portfolios. A value-weighted HML portfolio sorted on 90% rVaR has a monthly return of -0.42% with a t-statistic of -2.90. In contrast, a value-weighted HML portfolio sorted on 50% rVaR has a return of 0.66% with a t-statistic of 4.51. Moreover, the observed pricing effects persist after controlling for the Fama & French (2016) five factor model and the t-statistics even increase.

These findings are consistent with our explanation that investors prefer to rely on information from options close-to-the-money. Information obtained from traded options in the center of the distribution is reflected in the cross-section of returns. If the rVaR is estimated from options in the center, we observe that riskier stocks carry a premium consistent with downside risk-averse investors. In contrast, information obtained from the extrapolated tail is viewed as less reliable and is thus ignored by investors despite the information content in far-out-of-the-money options as shown by Chakravarty et al. (2004) or Augustin et al. (2015). Informed traders prefer to trade in out-of-the-money options as they offer greater leverage and thus incorporate information into far-out-of-the-money prices before it becomes apparent to investors trading in close-to-the-money options or the underlying stock. This combination of exploitable information in far-out-of-the-money options and the informational preference for close-to-the-money information likely generates the market inefficiency ob-

served by Rehman & Vilkov (2012), Schneider et al. (2015) and Stilger et al. (2016). Investors prefer to rely on information generated from liquid options which is believed to be more reliable. They base their decisions on option-implied information obtained from option data close-to-the-money. In contrast, far-out-of-the-money options are often illiquid and might even be based on an extrapolation of the volatility surface. Their informational content is therefore considered more unreliable.

Furthermore, we observe that the absolute return and t-statistic of the 95% rVaR sorted HML portfolio is smaller than the return and t-statistic of the 85% rVaR sorted portfolio. This finding is likely due to the declining quality of option data in the tail of the distribution. Sorting stocks on option-implied information based on the tail of the distribution requires that there is exploitable information in these options. If the underlying price data becomes more unreliable, the results become more noisy leading to lower t-statistics and lower absolute returns.

Our results also provide a potential explanation for previous contradictory results in the literature. Conrad et al. (2013) document that left-skewed stocks carry a premium over right-skewed stocks in the cross-section of returns. In contrast, Rehman & Vilkov (2012), Schneider et al. (2015) and Stilger et al. (2016) report that right-skewed stocks carry a premium. Conrad et al. (2013) use heavily filtered raw option data to estimate risk-neutral moments. This truncates the domain of available option prices, and their option-implied skewness estimates are strongly dependent on option prices close-to-the-money. Therefore, Conrad et al. (2013) seem to estimate skewness from the center and shoulder of the risk-neutral distribution and not from its tails. Consequently, their dataset is unlikely to include far out-of-the-money options and it appears plausible that Conrad et al. (2013) do not pick up the effect of skewness in the tails.

Moreover, our results of 90% rVaR are consistent with Rehman & Vilkov (2012)

Table 2: Portfolios sorted on option-implied rescaled Value-at-Risk

This table presents the returns of portfolios sorted on option-implied rescaled Value-at-Risk (rVaR) for the S&P 500 constituents. At the end of each month stocks are sorted into tercile portfolios with breakpoints of 0.33 and 0.67. Portfolios are kept constant for the next month. Option-implied rVaR is defined as $-\frac{q(1-p)}{q(0.75)-q(0.25)}$ where q is the quantile function. We begin by sorting stocks into buckets based on a value of $p = 0.95$. The results for the full sample are reported in column one. Then we repeat the procedure for $p = 0.90$ and report the results in column two. The estimation procedure is continued for values down to $p = 0.50$ (the median) in steps of 0.05. The results in each column are therefore independent of each other. Low, mid and high portfolios contain stocks in the first, second and third tercile. High-minus-low (HML) is the difference between the high and low portfolio. We report the equal-weighted, the value weighted and the Fama & French (2016) 5 factor model adjusted returns. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

rVaR	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Equal weighted										
Low	1.23	1.21	1.24	1.02	0.86	0.84	0.78	0.72	0.70	0.70
Mid	0.88	0.90	0.93	1.10	0.99	0.99	1.02	0.98	0.96	0.94
High	0.86	0.85	0.80	0.84	1.12	1.14	1.16	1.26	1.30	1.31
HML	-0.37	-0.36	-0.44	-0.18	0.26	0.30	0.38	0.54	0.61	0.61
t-Stat	-2.37	-2.55	-3.59	-1.12	1.38	2.36	2.90	4.14	4.35	3.95
Value weighted										
Low	1.04	1.03	1.01	0.76	0.74	0.70	0.64	0.46	0.39	0.43
Mid	0.71	0.74	0.79	0.88	0.81	0.84	0.86	0.83	0.84	0.85
High	0.64	0.61	0.57	0.74	0.75	0.81	0.81	1.08	1.13	1.09
HML	-0.40	-0.42	-0.44	-0.03	0.01	0.12	0.17	0.62	0.74	0.66
t-Stat	-2.67	-2.90	-2.95	-0.13	0.03	0.93	1.25	4.05	4.54	4.51
FF5 adjusted										
Low	0.24	0.21	0.20	-0.03	-0.14	-0.13	-0.16	-0.24	-0.27	-0.28
Mid	-0.10	-0.08	-0.03	0.11	-0.06	-0.02	-0.01	-0.01	-0.04	-0.06
High	-0.12	-0.13	-0.16	-0.07	0.21	0.16	0.18	0.26	0.32	0.35
HML	-0.36	-0.34	-0.37	-0.04	0.35	0.28	0.35	0.50	0.59	0.62
t-Stat	-2.97	-2.92	-3.54	-0.46	2.52	3.09	3.68	4.62	4.92	4.86

and Stilger et al. (2016) for rVaR with confidence levels larger than 80%. The traditional central skewness measure places a large weight on changes in probability mass far in the tail of the distribution. Pearson skewness is defined as $E\left[\frac{X-\mu}{\sigma}\right]^3$. The X^3 term causes the traditional skewness measure to be very sensitive to changes in probability mass in the tails. Rehman & Vilkov (2012) and Stilger et al. (2016) interpolate and extrapolate the implied volatility surface and thus their skewness measure is strongly dependent on the information contained in far-out-of-the-money options. The differences in results between Conrad et al. (2013), Rehman & Vilkov (2012) and Stilger et al. (2016) are consistent with our findings.

1.4.2 Competing information

Our proposed explanation for the occurrence of two distinct pricing effects of option-implied information implies that the observed effects should amplify if the information from the tail and center of the distribution are mismatching. If investors indeed prefer information obtained from close-to-the-money options over information from far-out-of-the-money options, contradictory information from the tails and center should lead to an even stronger pricing effect. We test this idea by performing two double sorts on 50% rVaR and 90% rVaR. Sorts are performed at the end of the last trading day every month and portfolios are kept constant in the subsequent month. The results are presented in Table 3.

Panel A presents the results of first sorting on information from close-to-the-money options (50% rVaR) and then sorting on information from far-out-of-the-money options (90% rVaR). More importantly, we also present the returns of the two diagonal portfolios. The High-Low minus Low-High portfolio (HL-LH) trades stocks which have contradictory information implied by the 50% rVaR and 90% rVaR. The portfolio is long in stocks which have a high 50% rVaR , indicating high risk, and a low 90% rVaR, indicating low risk (top right portfolio). It is short in stocks with low 50%

rVaR and high 90% rVaR (bottom left portfolio). The return of the HL-LH portfolio is 1.14% per month with a t-statistic of 4.83. This finding confirms our hypothesis that the pricing effects are strongest if the two signals are contradictory. If the information from close-to-the-money and far-out-of-the-money options disagree, investors prefer the reliable information source and ignore the information from the tail of the risk-neutral distribution.

We also create a High-High minus Low-Low (HH-LL) portfolio. The HH-LL portfolio is formed by taking a long position in the High-High portfolio (bottom right) containing stocks which have a high 50% rVaR and a high 90% rVaR and shorting the diagonally opposite Low-Low portfolio (top left) containing stocks which have a low 50% rVaR and a low 90% rVaR. The HH-LL portfolio yields a monthly return of 0.35 with a t-statistic of 1.99 which is in line with downside risk-aversion. Furthermore, we can observe that the 90% rVaR anomaly is strongest for stocks with a high 50% rVaR. This is likely due to more liquid option markets for stocks with a high option-implied 50% rVaR. If close-to-the-money options indicate a higher risk, more market participants buy insurance against downside moves providing additional liquidity thus making the 90% rVaR estimates less noisy.

Panel B shows the results of an inverted sort order. In a first step, we sort on the option-implied 90% rVaR, and in a second step we sort on the 50% rVaR . The results are comparable to the observations from Panel A and further confirm our hypothesis that competing information leads to an even stronger pricing effect. The HL-LH portfolio yields a monthly return of -0.90% with a t-statistic of -3.33.

1.4.3 The effect of liquidity

The previously observed pricing effect of 90% rVaR requires that there is information in the out-of-the-money put options that can be used to predict the cross-section

Table 3: Competing information

This table presents results of double sorted portfolios. In Panel A (B) stocks are first sorted into tercile portfolios at the end of each month based on the 30 day 50% (90%) rVaR and second on the 30 day 90% (50%) rVaR. Portfolios are kept constant for the next month. The table also displays the returns of the diagonal portfolios High-High minus Low-Low (HH-LL) and High-Low minus Low-High (HL-LH). Monthly return data covers the period from February 1996 to August 2015, with a total of 235 monthly observations for all S&P 500 constituents. Returns are value-weighted. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

		Panel A: 50% rVaR			Diagonal Portfolios	
		Low	Mid	High	HH-LL	HL-LH
90% rVaR	Low	0.51	0.87	1.51		
	Mid	0.48	0.91	1.08		
	High	0.36	0.79	0.86		
	HML	-0.15	-0.08	-0.65	0.35	1.14
	t-Stat	-0.73	-0.45	-2.90	1.99	4.83
		Panel B: 90% rVaR			Diagonal Portfolios	
		Low	Mid	High	HH-LL	HL-LH
50% rVaR	Low	0.81	0.44	0.36		
	Mid	1.05	0.67	0.68		
	High	1.26	1.15	0.79		
	HML	0.45	0.71	0.43	-0.02	-0.90
	t-Stat	2.28	3.38	2.21	-0.10	-3.33

of returns. If there is no additional information in the prices of far-out-of-the-money put options, e.g. because they are too illiquid, no information rent can be gained by sorting stocks on 90% rVaR. We confirm this intuition by performing a double sort on three proxies for option market liquidity: First, the liquidity of the underlying stock measured by Amihud (2002) illiquidity over the past month. The liquidity of the underlying is strongly related the liquidity of the option market (e.g. Cao & Wei, 2010). Second, the cumulative dollar volume of all put option contracts aggregated over the past month. Third, the cumulative open interest of all put options on the last trading day of the month. The results are presented in Table 4. Panel A presents the results for 50% rVaR. The results show that the predictive power of 50% rVaR does not interact with any liquidity proxies. For example, the double sort on put volume and 50% rVaR reveals that the HML portfolio sorted on 50% rVaR for stocks within the lowest 33% of put option volume has a return (t-statistic) of 0.53% (5.14). The return and t-statistic do not change significantly for stocks in the mid and upper third subsets of put volume where the HML portfolios have a return (t-statistic) of 0.66% (3.93) and 0.57% (3.10) respectively. The same finding can be observed if the first sort is performed on Amihud (2002) illiquidity or put open interest. This finding is in-line with our expectations. 50% rVaR is based on options close-to-the-money which are reasonably liquid within our sample of S&P500 stocks and therefore always contain enough extractable information.

The results in Panel B show the results of double sorts for 90% rVaR. We observe that the observed market inefficiency is strongest for stocks with relatively liquid option markets and disappears for stocks with illiquid option markets. Double sorting on put option volume reveals that the return predictability is driven by those stocks with liquid option markets. The HML portfolio for those stocks with a high put option volume yields a return of -0.42% per month with a t-statistic of -2.25, compared to an insignificant return of -0.20% for stocks with a low put option volume. The results are comparable among the three proxies for option market liquidity and

confirm that option-implied information obtained from far-out-of-the-money options can only be traded successfully if the option market is liquid enough.

1.4.4 Interactions with risk-neutral moments

We are interested how our results based on the rVaR measure compare to other studies that rely on the option-implied moments of Bakshi et al. (2003) such as Rehman & Vilkov (2012) and Stilger et al. (2016). Our interest is sparked by the observation that the cross-sectional pricing of 90% rVaR is similar to the findings of Rehman & Vilkov (2012), Stilger et al. (2016), and Schneider et al. (2015). In contrast, the results of 50% rVaR are comparable to those in Conrad et al. (2013). We expect that our results of 90% rVaR are more closely related to the results of Rehman & Vilkov (2012) and Stilger et al. (2016) because both studies extrapolate the implied volatility surface. Skewness and kurtosis are very sensitive to changes in the tail of the probability distribution as they include a power of 3 and 4 respectively. Thus, any option-implied skewness or kurtosis measure relies strongly on the information contained in far-out-of-the-money options. In contrast, Conrad et al. (2013) filter out all options with a bid below 50 cent and do not extrapolate the implied volatility surface. Therefore, their skewness and kurtosis measures are much more dependent on options which are relatively close-to-the-money which explains why their results are comparable to the findings of our 50% rVaR measure.

To further study the relation between the option-implied moments of Bakshi et al. (2003) and rVaR, we perform double sorts by first sorting on option-implied moments and then on rVaR. The option-implied moments of Bakshi et al. (2003) are estimated from the same volatility surface as our rVaR measures which is extrapolated horizontally and interpolated with a cubic spline similar to Carr & Wu (2009) and Rehman & Vilkov (2012). The results of the double sorts are presented in Table 5. Panel A shows the results of 50% rVaR double sorts. Our 50% rVaR measure is not related to

Table 4: The effect of liquidity

This table presents results of double sorted portfolios. At the end of each month stocks are first sorted on the liquidity of their markets. More specifically, we sort on Amihud (2002) illiquidity of the underlying stock in the past month (columns 1–3), dollar volume of all traded put options over the last month, and cumulative open interest (columns 7–9) of all traded put options on the last trading day. Within each tercile portfolio stocks are then sorted on their 50% rVaR (Panel A) or 90% rVaR (Panel B). Portfolios are kept constant for the next month. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations for all S&P 500 constituents. Returns are value-weighted. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

Panel A: 50% rVaR

		Illiquidty			**Put volume**			**Put open interest**		
		Low	Mid	High	Low	Mid	High	Low	Mid	High
50% rVaR	Low	0.42	0.75	0.73	0.66	0.63	0.46	0.75	0.56	0.44
	Mid	0.77	0.96	0.94	0.98	0.84	0.77	0.94	0.96	0.73
	High	0.97	1.20	1.23	1.20	1.29	1.03	1.19	1.32	1.08
	HML	0.55	0.45	0.50	0.53	0.66	0.57	0.43	0.76	0.63
	t-Stat	3.17	2.82	3.42	5.14	3.93	3.10	3.47	4.60	3.42

Panel B: 90% rVaR

		Illiquidty			**Put volume**			**Put open interest**		
		Low	Mid	High	Low	Mid	High	Low	Mid	High
90% rVaR	Low	1.00	1.09	1.07	1.08	1.14	0.97	1.06	1.09	0.95
	Mid	0.69	0.94	0.85	0.87	0.95	0.73	0.82	0.94	0.78
	High	0.53	0.87	0.96	0.88	0.68	0.55	0.90	0.64	0.54
	HML	-0.47	-0.22	-0.12	-0.20	-0.46	-0.42	-0.16	-0.45	-0.41
	t-Stat	-3.06	-1.78	-0.70	-1.79	-3.07	-2.25	-1.33	-3.14	-2.14

option-implied skewness and kurtosis of Bakshi et al. (2003). For example, the return (t-statistic) of the HML portfolio sorted on 50% rVaR for stocks in the low kurtosis subset is 0.61% (2.42). Among high kurtosis stocks the return (t-statistic) of the HML portfolio is 0.66% (2.86). The results are slightly weaker for option-implied skewness indicating that option-implied skewness captures some of the information contained in the 50% rVaR measure. Indeed, the difference between median and mean is sometimes used as a proxy for skewness. The mean of the risk-neutral distribution is the same for all stocks and thus the measures based on the median of the risk-neutral distribution may be related to skewness. The return (t-statistic) of the HML portfolio sorted on 50% rVaR for stocks in the low skewness subset is 0.40% (2.63). In the high skewness subset the return (t-statistic) of the HML portfolio is 0.39% (1.91). The results show that 50% rVaR is not strongly related to the option-implied moments of Bakshi et al. (2003) and the pricing of the center of the risk-neutral distribution thus constitutes a separate asset pricing phenomenon.

In contrast, 90% rVaR is strongly related to option-implied skewness and kurtosis. Panel B presents the results for 90% rVaR double sorts. Option-implied kurtosis completely subsumes 90% rVaR. The returns of the HML portfolios which are first sorted on option-implied kurtosis are -0.13%, 0.07%, and 0.24% for stock in the low, mid, and high kurtosis subsets respectively. The results are similar when the first sort is performed on option-implied skewness. The effect vanishes for the low and mid skewness subset. However, for stocks with relatively high skewness the 90% rVaR anomaly persists and even becomes stronger yielding a return (t-statistic) of -0.51% (-2.40). The results confirm our intuition that 90% rVaR is closely related to skewness and that it is a measure of tail risk, therefore providing further evidence that Rehman & Vilkov (2012) and Stilger et al. (2016) pick up the pricing effect from far-out-of-the-money put options. In contrast, Conrad et al. (2013) use the information of close-to-the-money options.

Table 5: Interaction with risk-neutral moments

This table presents results of double sorted portfolios. At the end of each month stocks are first sorted on the option-implied moments of Bakshi et al. (2003). More specifically, we sort on option-implied skewness (columns 1–3) and option-implied kurtosis (columns 4–6) . Within each tercile portfolio stocks are then sorted on their 50% rVaR (Panel A) or 90% rVaR (Panel B). Portfolios are kept constant for the next month. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations for all S&P 500 constituents. Returns are value-weighted. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

Panel A: 50% rVaR

		Impl. skewness			**Impl. kurtosis**		
		Low	Mid	High	Low	Mid	High
50% rVaR	Low	0.34	0.66	0.88	0.79	0.58	0.31
	Mid	0.50	0.79	1.07	1.17	0.73	0.65
	High	0.73	1.14	1.28	1.40	0.91	0.96
	HML	0.40	0.48	0.39	0.61	0.34	0.66
	t-Stat	2.63	1.89	1.91	2.42	1.93	2.86

Panel B: 90% rVaR

		Impl. skewness			**Impl. kurtosis**		
		Low	Mid	High	Low	Mid	High
90% rVaR	Low	0.66	0.96	1.42	1.13	0.65	0.61
	Mid	0.53	0.84	1.04	1.31	0.84	0.70
	High	0.51	0.87	0.91	0.99	0.72	0.38
	HML	-0.15	-0.09	-0.51	-0.13	0.07	-0.24
	t-Stat	-0.80	-0.53	-2.40	-0.45	0.32	-1.24

1.4.5 Interactions with low-risk anomalies

Schneider et al. (2015) link beta- and volatility-based low-risk anomalies, such as the low-beta anomaly (Frazzini & Pedersen, 2014), to option-implied skewness. They argue that standard asset pricing models fail for stocks with high downside risk as they cannot capture the non-linear relation between market returns and stock returns induced by downside risk. Therefore, these anomalies yield a skewness risk premium. We are therefore interested in the relation between our results, realized beta, and volatility. We estimate realized volatility as a rolling 1 year standard deviation computed from daily returns and beta from a rolling covariance matrix over a 5 year window and a 1 year standard deviation of market returns as in Frazzini & Pedersen (2014). The results are presented in Table 6.

Panel A shows the results for double sorts of 50% rVaR. The results show that the 50% rVaR pricing is stronger for mid and high levels of realized volatility than for low realized volatility stocks. HML portfolios of mid and high realized volatility stocks yield a monthly return (t-statistic) of 0.71% (3.40) and 0.54% (2.05) respectively. For low realized volatility the HML return (t-statistic) is 0.25% (1.53). A possible explanation for this finding is that the Low-Low portfolio generates a relatively high return of 0.63% per month which could be due to a low-risk anomaly. The results for the double sort of rolling beta and 50% rVaR show a similar pattern. The return of the HML portfolio is highest for mid and high beta stocks with returns (t-statistics) of 0.60% and (2.98) and 0.64% (3.03). In contrast, the return of the HML portfolio for the subset of low beta stocks is 0.49% (3.31). In conclusion, the 50% rVaR pricing seems to be largely unaffected by low-risk anomalies, except for the low risk subsets where the low-risk anomalies partially offset the pricing of 50% rVaR.

The results for 90% rVaR can be found in Panel B. In contrast to 50% rVaR, 90% rVaR interacts more strongly with realized volatility and beta. The anomaly is particularly strong among low risk stocks which is consistent with the findings of

Schneider et al. (2015). For example, the return (t-statistic) of the HML portfolio is -0.46% (-3.21) for low beta stocks. In contrast, the HML portfolio for the subset of high beta stocks has a return (t-statistic) of -0.32% (-1.40). A similar observation can be made for realized volatility. The effect is the most pronounced for low realized volatility stocks and insignificant for mid and high beta stocks. The results show that 90% rVaR is stronger among low risk stocks. However, both 50% and 90% rVaR persist when controlling for low-risk anomalies and are thus distinct effects.

1.4.6 Persistence of the return predictability

Conrad et al. (2013) and Rehman & Vilkov (2012) report that the returns from portfolios sorted on risk-neutral skewness are persistent for subsequent months. Therefore, we hypothesize that the return of portfolios sorted on rVaR is also persistent. The persistence of returns allows us to draw further conclusions about the source of this asset pricing phenomenon. An et al. (2014) argue that the return predictability in their study is caused by informed traders who prefer to trade in option markets, thus incorporating information in the option market first. Their model is similar to the model of Easley, O'Hara, & Srinivas (1998). If our results are indeed driven by informed traders, the return predictability should only persist in the first few days after portfolio formation and should disappear or even reverse afterwards. In contrast to a micro-structural informed trading explanation, option traders could be able to identify some underlying fundamental characteristics. Rehman & Vilkov (2012) and Stilger et al. (2016) show that the return predictability of risk-neutral skewness is particularly strong among overvalued firms.

In Table 7 we report the returns of portfolios sorted on rVaR 2, 3, 6, and 12 months after their initial portfolio formation. The reported returns are monthly returns in the respective months, not cumulative returns from the formation date. We also report the return for the first 5 trading days after portfolio formation and the

Table 6: Interaction with low-risk anomalies

This table presents results of double sorted portfolios. At the end of each month stocks are first sorted on other asset pricing phenomena. More specifically, we first sort 1 year rolling realized volatility (columns 1–3), and rolling beta as in Frazzini & Pedersen (2014) (columns 4–6) . Within each tercile portfolio stocks are then sorted on their 50% rVaR (Panel A) or 90% rVaR (Panel B). Portfolios are kept constant for the next month. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations for all S&P 500 constituents. Returns are value-weighted. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

Panel A: 50% rVaR

		Realized volatility			**Rolling beta**		
		Low	Mid	High	Low	Mid	High
50% rVaR	Low	0.63	0.49	0.59	0.54	0.40	0.42
	Mid	0.75	0.93	0.87	0.95	0.78	0.74
	High	0.88	1.20	1.14	1.03	1.00	1.06
	HML	0.25	0.71	0.54	0.49	0.60	0.64
	t-Stat	1.53	3.40	2.05	3.31	2.98	3.03

Panel B: 90% rVaR

		Realized volatility			**Rolling beta**		
		Low	Mid	High	Low	Mid	High
90% rVaR	Low	0.96	1.03	1.03	1.11	0.95	0.94
	Mid	0.66	0.68	1.02	0.77	0.72	0.73
	High	0.65	0.84	0.64	0.66	0.56	0.62
	HML	-0.31	-0.18	-0.39	-0.46	-0.39	-0.32
	t-Stat	-2.41	-0.84	-1.46	-3.21	-2.13	-1.40

return from the 6^{th} trading day until the end of the first month post formation.

The results show that the return predictability for rVaR measured far in the tail (85-95% rVaR) is persistent for up to 6 months after portfolio formation. For example, for 90% rVaR the return of the HML portfolio in the third month after portfolio formation is -0.39% with a t-statistic of -2.25. For subsequent months the average HML return remains significant up to 6 months after portfolio formation. Furthermore, the return in the first 5 days after portfolio formation is not statistically significant if rVaR is computed from far-out-of-the-money options and the effect sets in during the rest of the month. The return of the 90% rVaR HML portfolio generates a return (t-statistic) of -0.08% (-1.11) over the first 5 days whereas the return (t-statistic) over the subsequent trading days is -0.36 (-2.49). This observation is in line with the results of Rehman & Vilkov (2012) and Stilger et al. (2016). Both studies link the anomaly to over- or undervaluation of the underlying firms.

In contrast to the pricing of information from far-out-of-the-money options, option-implied information from options close-to-the-money carries more information for returns over a short time horizon. For example, returns of the HML portfolio sorted on 50% rVaR yield a return (t-statistic) of -0.03% (-0.18) and 0.18% (1.29) 2 and 3 months after portfolio inception respectively. In fact, most of the effect can be attributed to the first 5 trading days after the portfolio creation. The return (t-statistic) of the 50% rVaR HML portfolio during the first 5 days is 0.39% (4.18). In contrast, the return (t-statistic) during the remainder of the month drops to 0.28% (2.24) even though the time horizon is much longer. This phenomenon is similar to the findings of Goncalves-Pinto, Grundy, Hameed, van der Heijden, & Zhu (2016) who construct a measure based on the difference between option-implied stock price and traded price. Their measure relies mostly on options close-to-the-money. They find that the return predictability of their measure is strongest one day after portfolio formation. Goncalves-Pinto et al. (2016) conclude that their finding is likely related to

price pressure in the stock market. In conclusion, the results of Table 7 further show that the return predictability from far-out-of-the-money options is distinct from the return predictability of close-to-the-money options. The pricing effect of 50% rVaR is particularly present over short horizons, whereas the predictability derived from 90% rVaR persistent for up to 6 months.

1.5 Robustness

Measures of option-implied information are difficult to construct reliably. Therefore, we perform extensive tests to ensure that our results are not driven by our particular choice of methodology. First, we investigate if the effect is driven by artifacts in the data which could create the observed effects. Second, we test if our results persist if we choose a different interpolate and extrapolation methodology. Third, we scale our option-implied rVaR measure with the square root of option-implied variance. The robustness checks show that our results are robust to our choice of methodology.

1.5.1 Correlations between portfolio positions

Our results show that sorting returns on 50% rVaR and 90% rVaR yields results with opposing signs. The two effects occur over different time horizons and interact differently with other asset pricing phenomena. However, we want to rule out that our results are driven by artifacts in the data that "mechanically" create the change in sign between 50% and 90% rVaR. To ensure that the sorts of 90% rVaR are not merely inverted sorts of 50% rVaR, we compute Spearman's rank correlation between the portfolio positions of the stocks. A stock sorted into the upper tercile is assigned a portfolio position of 1, in the mid tercile 0, and in the lowest tercile -1. If 50% rVaR and 90% rVaR indeed measure the same effect with different signs, the portfolio positions should be inverted and Spearman's rank correlation should be close to -1.

Table 7: Persistence of option-implied rVaR

This table presents the returns of portfolios sorted on option-implied rescaled Value-at-Risk (rVaR) for the S&P 500 constituents. At the end of each month stocks are sorted into tercile portfolios with breakpoints of 0.33 and 0.67. Portfolios are kept constant for the next month. Option-implied rVaR is defined as $-\frac{q(1-p)}{q(0.75)-q(0.25)}$ where q is the quantile function. We begin by sorting stocks into buckets based on a value of $p = 0.95$. The results for the full sample are reported in column one. Then we repeat the procedure for $p = 0.90$ and report the results in column two. The estimation procedure is continued for values down to $p = 0.50$ (the median) in steps of 0.05. Returns are the returns during the respective period. Low, mid and high portfolios contain stocks in the first, second and third tercile. High-minus-low (HML) is the difference between the high and low portfolio. We report the equal-weighted, the value weighted and the Fama & French (2016) 5 factor model adjusted returns. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations. Returns are value-weighted. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

rVaR	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Return over the first 5 days										
Low	0.44	0.41	0.38	0.38	0.35	0.32	0.31	0.21	0.18	0.21
Mid	0.39	0.41	0.44	0.38	0.43	0.49	0.45	0.41	0.41	0.41
High	0.33	0.33	0.33	0.43	0.37	0.33	0.43	0.60	0.63	0.60
HML	-0.11	-0.08	-0.05	0.04	0.02	0.02	0.11	0.40	0.45	0.39
t-Stat	-1.29	-1.11	-0.74	0.45	0.29	0.23	1.59	4.19	4.84	4.18
Return from day 6 until the end of the month										
Low	0.65	0.66	0.67	0.40	0.41	0.40	0.35	0.27	0.23	0.24
Mid	0.33	0.35	0.37	0.54	0.39	0.37	0.44	0.44	0.45	0.47
High	0.32	0.29	0.26	0.32	0.43	0.51	0.41	0.51	0.52	0.52
HML	-0.33	-0.36	-0.41	-0.07	0.01	0.12	0.06	0.24	0.30	0.28
t-Stat	-2.18	-2.49	-2.77	-0.45	0.08	1.10	0.55	1.81	2.21	2.24
Return 2 months ahead										
Low	0.90	0.91	0.87	0.77	0.75	0.72	0.60	0.70	0.69	0.76
Mid	0.81	0.79	0.81	0.84	0.82	0.79	0.93	0.82	0.84	0.83
High	0.61	0.63	0.66	0.75	0.77	0.91	0.87	0.84	0.81	0.74
HML	-0.29	-0.28	-0.21	-0.02	0.02	0.20	0.27	0.13	0.12	-0.03
t-Stat	-1.73	-1.74	-1.32	-0.10	0.08	1.13	1.94	1.02	0.83	-0.18

continued on the next page

Table 7 continued

rVaR	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Return 3 months ahead										
Low	0.98	1.00	0.93	0.70	0.75	0.77	0.69	0.65	0.64	0.65
Mid	0.83	0.81	0.90	0.83	0.82	0.76	0.85	0.91	0.89	0.87
High	0.62	0.61	0.57	0.88	0.84	0.81	0.78	0.70	0.75	0.83
HML	-0.36	-0.39	-0.36	0.18	0.09	0.03	0.09	0.04	0.11	0.18
t-Stat	-1.90	-2.25	-2.39	1.01	0.40	0.27	0.64	0.34	0.88	1.29
Return 6 months ahead										
Low	0.96	0.99	0.88	0.82	0.71	0.68	0.66	0.66	0.69	0.68
Mid	0.86	0.85	0.94	0.89	0.75	0.81	0.85	0.84	0.81	0.80
High	0.56	0.54	0.57	0.69	0.97	0.92	0.85	0.86	0.86	0.88
HML	-0.40	-0.45	-0.31	-0.13	0.26	0.25	0.19	0.20	0.17	0.20
t-Stat	-2.50	-2.74	-2.16	-0.62	1.61	2.00	1.26	1.50	1.26	1.28
Return 12 months ahead										
Low	0.77	0.79	0.81	0.76	0.73	0.80	0.74	0.65	0.60	0.57
Mid	0.71	0.72	0.74	0.67	0.80	0.69	0.72	0.82	0.84	0.85
High	0.72	0.71	0.68	0.74	0.63	0.66	0.71	0.70	0.78	0.83
HML	-0.05	-0.08	-0.13	-0.02	-0.10	-0.14	-0.03	0.05	0.19	0.26
t-Stat	-0.33	-0.51	-0.76	-0.09	-0.62	-1.00	-0.22	0.40	1.51	1.62

Table 8: Spearman's Rank Correlation of Portfolio Positions

This table displays the Spearman's rank correlation of portfolio positions of stocks sorted on different rescaled Value-at-Risk (rVaR) for the S&P 500 constituents. At the end of each month stocks are sorted into tercile portfolios with breakpoints of 0.33 and 0.67. Portfolios are kept constant for the next month. Option-implied rVaR is defined as $-\frac{q(1-p)}{q(0.75)-q(0.25)}$ where q is the quantile function. We begin by sorting stocks into buckets based on a value of $p = 0.95$. We record the portfolio positions of stocks as 1, 0, and -1 for stocks sorted into the upper, mid, and lowest tercile respectively. Then we repeat the procedure for $p = 0.90$. The estimation procedure is continued for values down to $p = 0.50$ (the median) in steps of 0.05. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations. Spearman's rank correlations and rVaRs are reported in percent.

rVaR	50	55	60	65	70	75	80	85	90
55	80.14	100.00							
60	55.02	74.93	100.00						
65	23.48	38.43	63.06	100.00					
70	0.80	6.01	24.12	59.89	100.00				
75	-7.04	-12.81	-5.52	18.20	56.12	100.00			
80	-4.67	-9.97	-9.94	-5.93	7.03	35.00	100.00		
85	-18.53	-14.24	-11.64	-14.60	-22.27	-23.88	22.29	100.00	
90	-19.10	-11.20	-7.28	-11.53	-24.50	-35.32	2.73	66.92	100.00
95	-18.32	-9.55	-5.17	-9.19	-22.99	-36.10	-1.30	58.39	89.96

Table 8 shows Spearman's rank correlations of portfolio positions of stocks sorted on different rVaRs. The results show that the observed effects are distinct from each other. In particular, Spearman's rank correlation between the portfolio positions obtained from sorts on 50% and 90% rVaR is -19.10% which is relatively low. Moreover, we can recognize two clear clusters of portfolio positions. One cluster can be observed for portfolio positions based on 50-60% rVaR where Spearman's rank correlations are 55.02%, 74.93% and 80.14%. A second cluster can be observed for option-implied information based on far-out-of-the-money options. Spearman's rank correlations between portfolio positions sorted on 85-95% rVaR are 66.92%, 58.39%, and 89.96%. Therefore, the results support our finding that option-implied information is incorporated differently into the cross-section of returns, depending on whether it is obtained from far-out-of-the-money or close-to-the-money options.

1.5.2 A different interpolation and extrapolation methodology

To alleviate concerns that our results are driven by our particular choice of interpolation and extrapolation methodology, we repeat the estimation of risk-neutral moments with different estimators. In particular, we aim for a different extrapolation methodology of the implied-volatility smile. The most common alternative to horizontal extrapolation is linear extrapolation as described in G. J. Jiang & Tian (2007). Linear extrapolation is essentially a first order tailor approximation and has the advantage that it eliminates potential kinks at the last observed strike price. The slope at the bounds of the observed volatility surface is calculated with respect to the strike price. However, the main drawback of using linear extrapolation is that it adds noise to the estimates as the extrapolation is strongly dependent on the slope of the of the volatility smile at the boundary. Moreover, we follow the procedure of Aït-Sahalia & Lo (1998) who smooth the implied volatility surface with the local-constant or Nadaraya-Watson kernel estimator. However, in small samples the local-constant estimator tends to be inaccurate at the boundaries of the domain (e.g., Li & Racine, 2004). To avoid under- or over-estimation of the implied volatility at the boundary, we use the local-linear estimator proposed by Stone (1977) and Cleveland (1979). Similar to Aït-Sahalia & Lo (1998) we choose to use a Gaussian kernel and use leave-one-out cross-validation to select the bandwidth. The results of portfolio sorts on rVaR obtained from our alternative methodology are presented in Table 9.

The results show that our previous results are not dependent on our particular choice of methodology. The return (t-statistic) of an equal-weighted HML portfolio sorted on 95% rVaR is -0.33% (-3.41). In, contrast, the return (t-statistic) of a 55% rVaR HML portfolio is 0.31% (2.29). The results also persist after controlling for the five factors of Fama & French (2016). The adjusted return (t-statistic) of the 95% rVaR HML portfolio is -0.21% (-2.52) and 0.31 (2.87) for the 55% rVaR HML portfolio. However, results based on our alternative methodology are slightly weaker which is likely due to the increased noise of the linear extrapolation. Nevertheless,

the two distinct effects are still clearly observable and statistically significant.

1.5.3 Scaling with option-implied volatility

We choose to rescale option-implied VaR with the inter-quartile range. However, our results are unaffected if we instead scale with the square root of Bakshi et al. (2003) option-implied variance. The results are presented in Table 10.

The results are almost identical to the results in Table 2. The monthly return (t-statistic) of an equal-weighted HML portfolio sorted on 90% rVaR is -0.36% (-2.55) on average. Moreover, for 50% rVaR the return (t-statistic) is 0.61% (3.95). The results again become stronger if portfolios are value-weighted. Furthermore, the results persist after adjusting returns for the Fama & French (2016) five factor model.

1.6 Conclusion

We analyze the pricing of option-implied information based on a novel option-implied rescaled Value-at-Risk measure. We show that option-implied information is incorporated differently into the cross-section of returns depending on whether it is estimated from far-out-of-the-money or close-to-the-money options. If the rVaR is estimated from options close-to-the-money, stocks with higher implied risk outperform stocks with lower implied risk. For example, a HML portfolio sorted on 50% rVaR yields a monthly return (t-statistic) of 0.60% (4.51) per month. This finding is in line with downside-risk averse investors (e.g. Lettau et al., 2014). In contrast, if rVaR is estimated from far-out-of-the-money options, the pricing effect reverses and stocks with low implied risk outperform stocks with high implied risk. For example, a HML portfolio sorted on 90% rVaR yields a monthly return (t-statistic) of -0.42% (2.90) per month.

Table 9: Kernel regression with linear extrapolation

This table presents the returns of portfolios sorted on option-implied rescaled Value-at-Risk (rVaR) for the S&P 500 constituents. rVaR is estimated from a implied volatility surface smoothed with a kernel-regression and extrapolated linear in strike. At the end of each month stocks are sorted into tercile portfolios with breakpoints of 0.33 and 0.67. Portfolios are kept constant for the next month. Option-implied rVaR is defined as $-\frac{q(1-p)}{q(0.75)-q(0.25)}$ where q is the quantile function. We begin by sorting stocks into buckets based on a value of $p = 0.95$. The results for the full sample are reported in column one. Then we repeat the procedure for $p = 0.90$ and report the results in column two. The estimation procedure is continued for values down to $p = 0.50$ (the median) in steps of 0.05. The results in each column are therefore independent of each other. Low, mid and high portfolios contain stocks in the first, second and third tercile. High-minus-low (HML) is the difference between the high and low portfolio. We report the equal-weighted, the value weighted and the Fama & French (2016) 5 factor model adjusted returns. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

rVaR	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Equal weighted										
Low	1.15	1.11	0.94	0.87	0.83	0.83	0.87	0.83	0.87	0.85
Mid	0.99	0.97	1.05	1.04	1.00	0.97	0.89	0.96	0.92	1.01
High	0.82	0.88	0.96	1.05	1.13	1.16	1.21	1.17	1.18	1.10
HML	-0.33	-0.23	0.02	0.18	0.31	0.34	0.34	0.34	0.31	0.25
t-Stat	-3.41	-1.90	0.11	0.84	1.59	2.09	2.30	2.63	2.29	1.61
Value weighted										
Low	0.91	0.83	0.71	0.63	0.68	0.68	0.70	0.62	0.64	0.64
Mid	0.80	0.79	0.88	0.94	0.90	0.86	0.81	0.83	0.77	0.91
High	0.63	0.64	0.68	0.69	0.75	0.77	0.79	0.92	0.98	0.80
HML	-0.28	-0.19	-0.03	0.06	0.07	0.09	0.09	0.29	0.34	0.17
t-Stat	-1.69	-1.00	-0.14	0.29	0.38	0.69	0.72	2.47	2.83	0.97
FF5 adjusted										
Low	0.11	0.04	-0.11	-0.17	-0.17	-0.16	-0.09	-0.12	-0.10	-0.11
Mid	-0.00	-0.03	0.03	0.01	-0.04	-0.04	-0.13	-0.08	-0.09	-0.05
High	-0.10	-0.00	0.09	0.17	0.22	0.21	0.23	0.20	0.20	0.17
HML	-0.21	-0.05	0.20	0.34	0.40	0.36	0.33	0.32	0.31	0.29
t-Stat	-2.52	-0.47	1.49	2.13	2.80	2.98	2.97	3.33	2.87	2.53

Table 10: Scaling with option-implied volatility

This table presents the returns of portfolios sorted on option-implied Value-at-Risk rescaled with the implied volatility of Bakshi et al. (2003) for the S&P 500 constituents . At the end of each month stocks are sorted into tercile portfolios with breakpoints of 0.33 and 0.67. Portfolios are kept constant for the next month. Rescaled option-implied VaR is defined as $-\frac{q(1-p)}{\sigma}$ where q is the quantile function and σ is the square root of Bakshi et al. (2003) implied variance. We begin by sorting stocks into buckets based on a value of $p = 0.95$. The results for the full sample are reported in column one. Then we repeat the procedure for $p = 0.90$ and report the results in column two. The estimation procedure is continued for values down to $p = 0.50$ (the median) in steps of 0.05. The results in each column are therefore independent of each other. Low, mid and high portfolios contain stocks in the first, second and third tercile. High-minus-low (HML) is the difference between the high and low portfolio. We report the equal-weighted, the value weighted and the Fama & French (2016) 5 factor model adjusted returns. Monthly return data cover the period from February 1996 to August 2015, for a total of 235 monthly observations. Standard errors are Newey & West (1987) adjusted with five lags. All results, except t-statistics, are reported in percent.

rVaR	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Equal weighted										
Low	1.23	1.21	1.24	1.02	0.86	0.84	0.78	0.72	0.70	0.70
Mid	0.88	0.90	0.93	1.10	0.99	0.99	1.02	0.98	0.96	0.94
High	0.86	0.85	0.80	0.84	1.12	1.14	1.16	1.26	1.30	1.31
HML	-0.37	-0.36	-0.44	-0.18	0.26	0.30	0.38	0.54	0.61	0.61
t-Stat	-2.37	-2.55	-3.59	-1.12	1.38	2.36	2.90	4.14	4.35	3.95
Value weighted										
Low	1.04	1.03	1.01	0.76	0.74	0.70	0.64	0.46	0.39	0.43
Mid	0.71	0.74	0.79	0.88	0.81	0.84	0.86	0.83	0.84	0.85
High	0.64	0.61	0.57	0.74	0.75	0.81	0.81	1.08	1.13	1.09
HML	-0.40	-0.42	-0.44	-0.03	0.01	0.12	0.17	0.62	0.74	0.66
t-Stat	-2.67	-2.90	-2.95	-0.13	0.03	0.93	1.25	4.05	4.54	4.51
FF5 adjusted										
Low	0.24	0.21	0.20	-0.03	-0.14	-0.13	-0.16	-0.24	-0.27	-0.28
Mid	-0.10	-0.08	-0.03	0.11	-0.06	-0.02	-0.01	-0.01	-0.04	-0.06
High	-0.12	-0.13	-0.16	-0.07	0.21	0.16	0.18	0.26	0.32	0.35
HML	-0.36	-0.34	-0.37	-0.04	0.35	0.28	0.35	0.50	0.59	0.62
t-Stat	-2.97	-2.92	-3.54	-0.46	2.52	3.09	3.68	4.62	4.92	4.86

These findings are in line with investors who prefer information from reliable sources, such as options close-to-the-money, over information embedded in unreliable sources, such as options far-out-of-the-money. Information obtained from traded options in the center of the distribution is reflected in the cross-section of returns, and is consistent with downside risk-aversion. If the rVaR is estimated from options close-to-the-money, we observe that riskier stocks carry a premium consistent with downside risk-averse investors. In contrast, information obtained from the far-out-of-the-money options is viewed as less reliable and is thus neglected by investors. However, there is evidence that there is usable information in far-out-of-the-money options. Chakravarty et al. (2004) and Augustin et al. (2015) argue that informed traders prefer to trade in out-of-the-money options as they offer greater leverage. The combination of exploitable information in far-out-of-the-money options and the informational preference for close-to-the-money information generates the observed market anomaly. If stocks are sorted on information obtained from far-out-of-the-money options, stocks with lower implied risk outperform stocks with higher implied risk, which is not in line with economic intuition.

The observed pricing effect is especially strong if the information from the close-to-the-money and far-out-of-the-money options is contradictory. This observation is consistent with investors who prefer information obtained from close-to-the-money options over information from far-out-of-the-money options. In particular, we perform a double sort on 50% rVaR and 90% rVaR and find that the diagonal High-Low minus Low-High portfolio yields a monthly return (t-statistic) of 1.14 (4.83).

Moreover, our results provide an explanation for contradictory results on the pricing of option-implied skewness in previous studies. Investors are downside risk-averse and want to be compensated for exposure to expected downside tail-risk. Therefore, negatively skewed stocks should have a higher return in the cross-section. This hypothesis is confirmed by Conrad et al. (2013), who find that stocks with nega-

tive skewness carry a premium compensating investors for the asymmetry of returns. However, studies by Rehman & Vilkov (2012) and Stilger et al. (2016) present contradictory evidence showing that stocks with high positive skewness earn a premium. All studies rely on the option-implied moments of Bakshi et al. (2003) but implement them differently. The study of Conrad et al. (2013) relies on to a larger extend on traded and relatively liquid options. In contrast, the study of Rehman & Vilkov (2012) and Stilger et al. (2016) extrapolate the implied volatility surface and thus rely to a larger extent on the information content of far-out-of-the-money options. The findings of these studies are consistent with our results. Conrad et al. (2013) appear to measure skewness from close-to-the-money options and thus find an effect consistent with the pricing of 50% rVaR. In contrast, Rehman & Vilkov (2012) and Stilger et al. (2016) find an effect which is consistent with the pricing of 90% rVaR.

2 Robust Estimation of Risk-Neutral Moments

Manuel Ammann & Alexander Feser

Status: Accepted for publication in the Journal of Futures Markets

ABSTRACT

This study provides an in-depth analysis of how to estimate risk-neutral moments robustly. A simulation and an empirical study show that estimating risk- neutral moments presents a trade-off between (1) the bias of estimates caused by a limited strike price domain and (2) the variance of estimates induced by micro-structural noise. The best trade-off is offered by option-implied quantile moments estimated from a volatility surface interpolated with a local-linear kernel regression and extrapolated linearly. A similarly good trade-off is achieved by estimating regular central option-implied moments from a volatility surface interpolated with a cubic smoothing spline and flat extrapolation.

Keywords: risk-neutral moments, risk-neutral distribution

JEL classification: C14, G10, G13, G17.

We thank Yakov Amihud, Fred Benth, Magnus Dahlquist, Ralf Elsas, Joachim Grammig, Fabian Hollstein, Kjell Nyborg, Jens Jackwerth, Joël Peress, Marcel Prokopczuk, Stefan Ruenzi, Nic Schaub, Christian Schlag, Paul Söderlind, Erik Theissen, Grigory Vilkov, as well as the participants of the annual conference of the Swiss Society for Financial Market Research 2017, of the 23rd annual meeting of the German Finance Association 2016, of the doctoral workshop at the 23rd annual meeting of the German Finance Association 2016, the Brown Bag Seminar of the University of St.Gallen, the 2016 joint seminar session of the University of St.Gallen and the University of Konstanz, and the 2015 Topics in Finance Seminar in Davos for helpful discussions and comments. All errors are our own.

2.1 Introduction

Information embedded in option prices is valuable for practitioners, regulators and academics alike. Unsurprisingly, an extensive literature has developed around the use of option-implied information (e.g. Carr & Wu, 2009; Chang, Christoffersen, Jacobs, & Vainberg, 2011; Buss & Vilkov, 2012; Conrad et al., 2013; Kozhan et al., 2013; Jurek, 2014; Stilger et al., 2016). Yet, there is surprisingly little literature which rigorously examines the efficiency of different risk-neutral moment estimators. In theory, the estimation of risk-neutral moments based on the theorem of Breeden & Litzenberger (1978) is an easy exercise given that a continuum of option prices is available. However, the practical estimation of risk-neutral densities from empirical option data is subject to many biases from discrete option prices and micro-structural noise (e.g. Bliss & Panigirtzoglou, 2002; Dennis & Mayhew, 2009). These biases may also be reflected in the results of empirical studies. For example, the question of how option-implied skewness is priced cannot be answered conclusively. Conrad et al. (2013) find that right skewed stocks carry a lower return than left skewed stocks, whereas Rehman & Vilkov (2012) find the opposite. Both studies use the method of Bakshi et al. (2003) to obtain risk-neutral skewness, but implement it differently. Thus, the differences in results could be driven by the difference in the estimation method. The main two studies analyzing the efficiency of risk-neutral moment estmators, Dennis & Mayhew (2009) and Bliss & Panigirtzoglou (2002), rely on early implementations of risk-neutral moments and only provide comparisons between a few select estimators. For example, Dennis & Mayhew (2009) do not use any interpolation or extrapolation, whereas the study of Bliss & Panigirtzoglou (2002) was published before the seminal work of Bakshi et al. (2003) and does not incorporate the use of the central moment formulas of Bakshi et al. (2003). We fill this gap in the literature by providing a comprehensive overview over the efficiency of the central moments of Bakshi et al. (2003), quantile-moments, the simple VIX (SVIX) of Martin (2017), and the rare disaster index (RIX) of Gao, Gao, & Song (2018); Gao, Lu, & Song (2018).

This study makes two contributions to the literature. First, it analyzes the efficiency of different estimation techniques for the popular risk-neutral moments of Bakshi et al. (2003), the simple VIX (SVIX) of Martin (2017), and the rare disaster index (RIX) of Gao, Gao, & Song (2018); Gao, Lu, & Song (2018) in-depth. A Monte Carlo simulation is used to price options under a stochastic volatility and jump (SVJ) model (Bates, 1996). The SVJ model nests the Black & Scholes (1973), the Heston (1993), and the Merton (1976) model and generates a distribution with higher levels of skewness and kurtosis than the normal distribution. Based on this simulated option data a classical horse race between three popular risk-neutral moment estimation methods is performed. In the horse race risk-neutral moments estimated from an implied volatility surface obtained from a cubic smoothing spline with horizontal extrapolation (e.g. Carr & Wu, 2009; Rehman & Vilkov, 2012; Neumann & Skiadopoulos, 2013), from a cubic smoothing spline with linear extrapolation (e.g. G. J. Jiang & Tian, 2007), from a local-linear kernel regression with linear extrapolation (Song & Xiu, 2016), and from a local-constant kernel regression with linear extrapolation (Aït-Sahalia & Lo, 1998) are tested. The robustness of the estimation methods is tested under real-world data quality by restricting the range of available strike prices, increasing the spacing between strikes, and by adding noise to the price data. All methods deliver accurate results if a large number of options over a wide range is available. As the domain spanned by available option prices declines, methods that extrapolate the implied volatility surface linearly in strike still perform reasonably well whereas horizontal extrapolation of the volatility surface performs poorly. In contrast, horizontal extrapolation is less affected by high levels of micro-structural noise, whereas methods which rely on linear extrapolation are strongly affected by small levels of micro-structural noise. This finding represents a bias-variance trade-off in real-world datasets: Researchers can choose methods that minimize the estimation error caused by a limited strike price domain or they can choose to minimize the variance of estimates induced by micro-structural noise in option-prices.

The second contribution is to propose the use of quantile moments to describe the risk-neutral distribution. Quantile moments are more robust and allow for greater flexibility than central moments. In contrast to central moments, quantile moments do not rely on probability weighting of outcomes making them more robust to the choice of extrapolation method and to data errors in far out-of-the-money options. Under real-world data quality the tails of the risk-neutral distribution are rarely observed as they require valid prices of far out-of-the-money options. Therefore, researchers will typically have to extrapolate the implied volatility surface or the risk-neutral distribution to artificially obtain prices of far out-of-the options. The extrapolation of the volatility surface requires an implicit assumption on the shape of the tail of the risk-neutral distribution. Central moments place a high weight on the tails of the distribution. For example, central skewness probability weights the cubed return thus placing an over-proportional weight on the prices of far out-of-the-money options.[1] Therefore central moments react sensitively to the choice of the extrapolation method. In contrast, quantile moments are computed by comparing the position of quantiles and thus are less sensitive to the choice of extrapolation method. Hinkley's (1975) measure of quantile skewness is used in this study. Hinkley's skewness compares the distance between the median and a quantile in the right tail to the distance between the median and a quantile in the left tail. The choice of quantiles provides additional flexibility compared to central moments, allowing researchers to measure the symmetry of the distribution at different points. Quantile kurtosis is measured by Ruppert's (1987) ratio of quantile ranges and quantile volatility is defined as the inter-quartile range. Quantile moments are by construction more robust to sparsely available option prices over a narrow domain than central moments. This intuition is confirmed in our simulation study. The findings show that quantile moments deliver accurate estimates of risk-neutral moments even if strike prices are truncated to a small domain. However, despite their practical and theoretical advantages, quantile moments have only been used rarely in the literature, for example Mirkov, Pozdeev,

[1] There seems to be evidence that traders use the over-proportional weight of far out-of-the-money options to manipulate central moments (Griffin & Shams, 2017).

& Söderlind (2019) use Hinkley's skewness to measure uncertainty around the removal of Swiss Franc cap. To the best of our knowledge, this is the first paper to propose the use of quantile kurtosis and inter-quartile range to describe the shape of the risk-neutral distribution.

The estimation of risk-neutral moments from real option data is subject to many biases. Option prices that span only a small domain truncate the available information. Consequently, narrow strike price domains supply us only with accurate information close to the current stock price, whereas tail information is lost. Furthermore, if the difference between two adjacent strike prices is large, the information is more sparse leading to potentially inaccurate estimates of the risk-neutral distribution. In addition to the bias caused by discrete prices, micro-structural noise in option prices is another important source of variance in estimates of risk-neutral moments. Observed option prices in empirical datasets are noisy. Prices are usually reported at a daily frequency, but often the last trade in each option contract happened at a different time of the day. This asynchroneous trading is a source of micro-structural noise. In addition, option prices have a bid-ask spread which requires to make an assumption on the true value of the option. Typically, the true value of the option is assumed to be the mid-price, but it could be anywhere between bid- and ask-price (Bliss & Panigirtzoglou, 2002).

Our study is most closely related to Dennis & Mayhew (2009) and Bliss & Panigirtzoglou (2002). Dennis & Mayhew (2009) simulate a range of European option prices from the Black-Scholes model. Therefore, returns follow a normal distribution and the real values of the risk-neutral moments are known and can be compared to the estimates. The drawback of the approach of Dennis & Mayhew (2009) is that they can only evaluate the bias in option prices which are based on a normal distribution, which is inconsistent with a the commonly observed implied volatility skew. Moreover, they do not interpolate or extrapolate the volatility surface, which is the

approach used by most recent studies such as Conrad et al. (2013) or Jurek (2014). Nevertheless, their results show that too large gaps between option prices or an insufficiently small range of option prices causes errors in option-implied moments. In particular they increase the spacing between strikes from 0.1 to 5 dollar in intervals of 10 cents. They find that for a stock with a current price of 70 dollar and a volatility of 20% risk-neutral skewness and kurtosis will start to oscillate around their true values. The amplitude is initially small but induces non-neglectable errors at larger spacings. For example, at a spacing just above 4.5 dollar between strikes, risk-neutral skewness takes a value smaller than -0.4, compared to the true value of 0. Dennis & Mayhew (2009) make a similar finding about the domain width, i.e. how far option data extends into the tails. In their normally distributed example of a stock with a volatility of 20% the skewness estimates are only unbiased if the strike prices extend about 20% into both tails. Bliss & Panigirtzoglou (2002) test the stability of a lognormal mixture model similar to Söderlind & Svensson (1997) and a method based on cubic smoothing splines similar to Jackwerth & Rubinstein (1996). They use real option data and add normally distributed noise to the prices to simulate micro-structure noise. Their results indicate that the smoothing spline implementation similar to Jackwerth & Rubinstein (1996) is more robust to noise than mixture based methods. In particular they find that the mixture based model often leads to unstable solutions or produced spurious spikes. In contrast to our study, they do not analyze the bias induced by small domains and discrete option data on the estimates of risk-neutral moments.

The results in our study go far beyond the results in Dennis & Mayhew (2009) and Bliss & Panigirtzoglou (2002). Our findings show that the properties of option-implied moment estimators depend on the interplay between the inter- and extrapolation method. This has important and actionable consequences for researchers implementing option-implied moment estimators. For example, linear extrapolation combined with local-linear kernel regressions leads to less noise sensitive estimates

than linear extrapolation combined with cubic smoothing splines, while having a similar bias due to narrow domain widths. Researchers that are concerned with low bias estimates of risk-neutral moments and that have access to relatively clean option prices, such as foreign exchange options, should thus prefer linear extrapolation combined with local-linear kernel regressions. In contrast, if researchers are estimating option-implied information from relatively noisy options, such as single stock equity options, horizontal extrapolation combined with cubic smoothing splines should be preferred over linear extrapolation combined with local-linear kernel regressions.

The paper proceeds as follows. Section 2.2 briefly reviews the risk-neutral moments of Bakshi et al. (2003) and introduces quantile moments. Section 2.3 discusses the different estimation techniques. Section 2.4 presents the methodology and the result of the horse-race between the different estimators. Section 2.5 presents our empirical results. Section 2.6 presents alternative specifications of the simulation study. Section 2.7 concludes.

2.2 Estimating Risk-Neutral Moments

Most of the modern literature (e.g. Bakshi et al., 2003; Carr & Wu, 2009; Kozhan et al., 2013; Martin, 2017) on risk-neutral moments relies on the theorem of Breeden & Litzenberger (1978) to obtain the risk-neutral density. The Breeden-Litzenberger theorem does not make an assumption about the price process of stocks and therefore allows recovery of the probability density function in a model-free way as:

$$F(S_T < K) = e^{r\tau} \frac{\partial P}{\partial K} \tag{4}$$

$$f(S_T) = e^{r\tau} \frac{\partial^2 P}{\partial K^2} \tag{5}$$

where P is the price of a put, r is the risk-free rate, τ is the time to maturity, F is the cumulative density function (CDF), and f is the probability density function (PDF) of the underlying under the risk-neutral measure. Hence, the CDF or the

PDF of the option-implied price distribution can be obtained by estimating the first or second derivative of a put option with respect to the strike price.

2.2.1 The risk-neutral moments of Bakshi, Kapadia and Madan (2003)

Bakshi et al. (2003) develop an analytical solution to obtain estimates of risk-neutral moments without obtaining the PDF first. They construct three synthetic securities that pay the squared, cubic, and quartic return at maturity respectively. The result of Bakshi & Madan (2000) allows to derive the analytical value of these contracts as:

$$\begin{aligned} V[t,\tau] = &\int_{S_t}^{\infty} \frac{2(1-ln[\frac{K}{S_t}])}{K^2} C(t,\tau,K)dK \\ &+ \int_{0}^{S_t} \frac{2(1-ln[\frac{S_t}{K}])}{K^2} P(t,\tau,K)dK \end{aligned} \tag{6}$$

$$\begin{aligned} W[t,\tau] = &\int_{S_t}^{\infty} \frac{6ln[\frac{K}{S_t}] - 3(ln[\frac{K}{S_t}])^2}{K^2} C(t,\tau,K)dK \\ &- \int_{0}^{S_t} \frac{6ln[\frac{S_t}{K}] - 3(ln[\frac{S_t}{K}])^2}{K^2} P(t,\tau,K)dK \end{aligned} \tag{7}$$

$$\begin{aligned} X[t,\tau] = &\int_{S_t}^{\infty} \frac{12(ln[\frac{K}{S_t}])^2 - 4(ln[\frac{K}{S_t}])^3}{K^2} C(t,\tau,K)dK \\ &+ \int_{0}^{S_t} \frac{12(ln[\frac{S_t}{K}])^2 - 4(ln[\frac{S_t}{K}])^3}{K^2} P(t,\tau,K)dK \end{aligned} \tag{8}$$

The challenge in the estimation of central risk-neutral moments is to evaluate the integrals in Equations (6) - (8). A numerical approximation of the integrals only delivers a reasonable approximation if a densely spaced set of option prices is available over a wide domain. The need for a wide domain is especially pressing for higher order moments that place a large weight on the tails of the distribution and thus on far out-of-the-money options. Inserting the prices of these contracts into the definitions

of central moments yields:

$$Vol^{\mathbb{Q}} = \sqrt{e^{r\tau}V - \mu^2} \tag{9}$$

$$Skew^{\mathbb{Q}} = \frac{e^{r\tau}W - 3e^{r\tau}\mu V + 2\mu^3}{[e^{r\tau}V - \mu^2]^{3/2}} \tag{10}$$

$$Kurt^{\mathbb{Q}} = \frac{e^{r\tau}X - 4\mu W + 6e^{r\tau}\mu^2 V - \mu^4}{[e^{r\tau}V - \mu^2]^2} \tag{11}$$

where

$$\mu = e^{r\tau} - 1 - e^{r\tau}\frac{V}{2} - e^{r\tau}\frac{W}{6} - e^{r\tau}\frac{X}{24} \tag{12}$$

2.2.2 Quantile Moments

Quantile moments are a convenient alternative to describe the shape of probability distributions. The quantile function is the inverse of the CDF and thus maps from the interval $[0, 1]$ onto the real line. The quantile function is obtained by numerically inverting the estimated CDF. Quantile moments, such as the median, have the advantage that they are more robust in the presence of data errors and outliers than traditional central moments. The idea behind quantile moments is to compare the relative position of quantiles to describe the shape of a probability distribution. Thus, it is possible to describe the shape of different portions of the probability distribution depending on the choice of the quantiles. For example, quantile skewness can be used to measure the symmetry in the center or tails of the distribution, making quantile moments more flexible than central moments. The principle behind quantile moments is illustrated in Figure 1. Their main difference from central moments is that they are not based on probability weighting of outcomes. Therefore, quantile moments react less to extreme events in the tails, making them inherently more robust than central moments. The inter-quartile range is used to measure quantile volatility which is defined as:

$$\text{QVol} = Q(0.75) - Q(0.25) \tag{13}$$

Inter-quartile range is a measure of dispersion that corresponds to the standard deviation of a probability distribution. It can be interpreted as the expected 50% confidence interval of the stock return (Figure 1 b). If the quantiles are further apart, the uncertainty about future returns is higher. It is also possible to compute the inter-quantile range from different quantiles given that they have the same distance from the median.[2] Quantile skewness allows for an equal flexibility. Hinkley (1975) defines quantile skewness as:

$$\text{QSkew}(p) = \frac{[Q(p) - Q(0.5)] - [Q(0.5) - Q(1-p)]}{Q(p) - Q(1-p)} \tag{14}$$

where $0.5 < p < 1$. Quantile skewness can take values between -1 and 1 whereas a value of 0 indicates that the distribution is symmetric (Groeneveld & Meeden, 1984). Quantile skewness is the normalized difference in the distance of the p quantile to the median and the $1-p$ quantile to the median. Therefore, it compares the length of the right tail to the length of the left tail (Figure 1c). Economically, quantile skewness measures whether there is more (risk-neutral) upside potential than downside risk. The study relies on Ruppert's (1987) measure of quantile kurtosis which is defined as:

$$\text{QKurt}(p) = \frac{Q(p) - Q(1-p)}{Q(q) - Q(1-q)} \tag{15}$$

where $0.5 < q < p < 1$. Quantile kurtosis is the ratio between two quantile ranges. It is always positive and larger than 1. Quantile kurtosis measures how far the tails extend in comparison to a reference interval. We choose $p = 0.95$ and $q = 0.75$ because in this case we scale with the inter-quartile range. The flexibility in the choice of p and q is an advantage over the traditional central kurtosis. Central kurtosis is

[2] An interesting case arises if the confidence interval is increased to nearly 100% because then the confidence interval of the risk-neutral measure $\mathbb{Q}$ is identical to the confidence interval of the real probability measure $\mathbb{P}$. The idea behind risk-neutral pricing is that in a complete and arbitrage-free market the risk-neutral measure $\mathbb{Q}$ exists and is equivalent to the real probability measure $\mathbb{P}$. The definition of equivalence states that if $\mathbb{Q}$ is equivalent to $\mathbb{P}$, both probability measures must agree on the states of the world with zero probability. Therefore, the 0^{th} and 100^{th} percentiles of both distributions must be identical. Even though the information embedded in the confidence interval is less exhaustive than the information of the complete physical density, it can be obtained with mild assumptions (uniqueness of the risk-neutral density) compared to e.g., Ross (2015).

difficult to define as it measures tail-weight and peakedness at the same time. Balanda & MacGillivray (1988) define central kurtosis as the "location- and scale-free movement of probability mass from the shoulders of a distribution into its center and tails". Schmid & Trede (2003) discuss that depending on the choice of p and q, quantile kurtosis can measure either tail-weight or peakedness. For large values of p, quantile kurtosis measures tail-weight, but for small values it measures peakedness. Economically, quantile kurtosis can be interpreted as an indication of how large extreme returns are in comparison to normal returns.

Obtaining the quantile function requires an estimated CDF. The CDF is obtained using the theorem of Breeden & Litzenberger (1978) and estimate the derivatives with central differences[3].

2.2.3 A naïve quantile approximation

In addition to the model-free method based on the theorem of Breeden & Litzenberger (1978) this study tests a naïve approximation of the CDF based on the model of Black & Scholes (1973) and Merton (1976). The Black-Scholes model states that the price of a put is:

$$P(S, K, \sigma_K, \tau, r) = e^{-r\tau} K \Phi(-d_2(\sigma)) - S_0 \Phi(-d_1(\sigma)) \tag{16}$$

where Φ is the CDF of the normal distribution, σ is the implied volatility, and d_1 and d_2 are defined as usual. Differentiating once with respect to the strike and rearranging yields:

$$F(S_\tau < K) = e^{r\tau} \Phi(-d_2(\sigma)) \tag{17}$$

[3]The choice of numerical differentiation method is not of particular relevance as we take the derivative of the interpolated volatility surface which has a very dense continuum of strike prices. Using other differentiation methods, e.g. a five point approximation, does not change the results.

Figure 1: Quantile moments

This figure illustrates the principle behind quantile moments. The plotted distribution is a sinh-arcsinh transformed normal distribution (Jones & Pewsey, 2009).Quantile volatility is the inter-quartile range which is defined as the difference between the 75^{th} and 25^{th} percentile. Quantile skewness is defined by Hinkley's (1975) measure at the 90^{th} percentile. Quantile kurtosis is defined by Ruppert's (1987) ratio of quantile ranges with an outer range between the 95^{th} and 5^{th} percentile and an inner range between the 75^{th} and 25^{th} percentile. However, the quantiles can be varied to describe a different portion of the probability distribution.

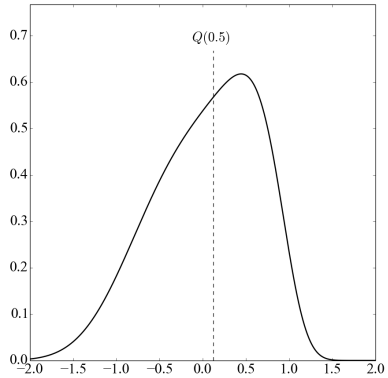

(a) Median

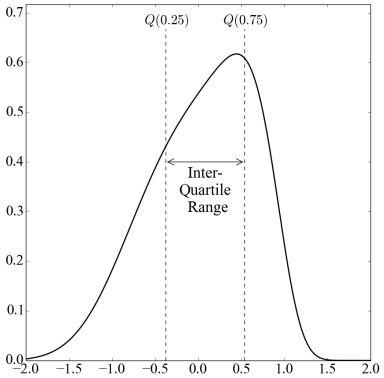

(b) Quantile volatility

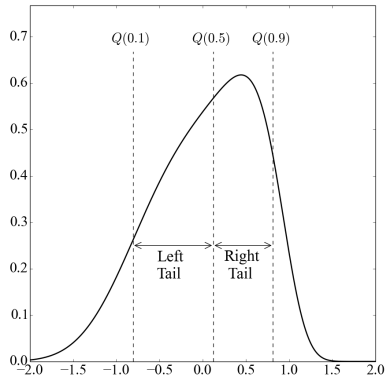

(c) Quantile skewness

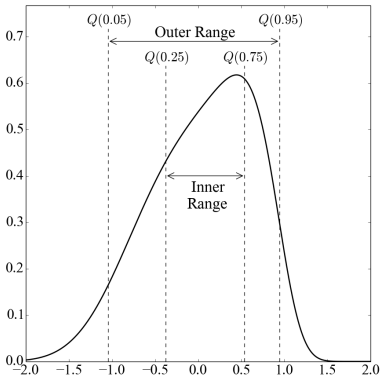

(d) Quantile kurtosis

Hence, the CDF of the stock price is naïvely recovered from the Black-Scholes formula. Note that the naïve approximation of the CDF is biased as it does not account for the implied volatility skew.[4] However, we opt to include the naïve approximation as it avoids taking a numerical derivative and should therefore be more stable than the model-free methodology.

2.2.4 Beyond option-implied moments: RIX, VIX, SVIX

The recent literature introduces a number of innovative option-implied measures. In particular, we analyze three measures in-depth. To begin with, Martin (2017) introduces a simplified volatility index SVIX which estimates option-implied variance more robustly. Traditional variance swaps and the VIX are based on the entropy of the underlying return process, i.e. they estimate the risk-neutral variance from log returns. However, the entropy based definition of the VIX only provides an unbiased estimate of volatility if the underlying process does not contain any jumps. Martin (2017) points out additional deficiencies of the traditional variance swaps, e.g. if the underlying stock goes bankrupt, the payoff of a variance swap would be infinite. To alleviate these problems, Martin (2017) proposes to estimate the variance from simple returns:

$$SVIX_t = \frac{2e^{r\tau}}{\tau F_t(\tau)^2}\left(\int_0^{F_t(\tau)} P(t,\tau,K)dK + \int_{F_t(\tau)}^{\infty} C(t,\tau,K)dK\right) \tag{18}$$

where $F_t(\tau)$ is the forward price at time t with maturity τ. The major difference of the SVIX to the VIX is that the VIX measures entropy whereas the SVIX measures variance. Entropy is more sensitive to the left tail of the distribution and thus the VIX loads more strongly on out-of-the-money put options. It is possible to construct a measure of non-lognormality by comparing VIX and SVIX (Martin, 2017). This idea is also used by Gao, Gao, & Song (2018) and Gao, Lu, & Song (2018) who

[4] Practitioners often use a skew correction (Gatheral, 2011, p. 104) which involves the derivative of the implied volatility surface with respect to the strike price and is essentially identical to the model-free methodology.

propose a rare disaster concern index (RIX) incorporating all higher-order moments. Their RIX measure is essentially the left-tail difference between the VIX and the jump-invariant version of the VIX similar to the variance measure of Bakshi et al. (2003) in Equation 6:

$$RIX = \frac{2e^{r\tau}}{\tau}\left(\int_0^{S_t} \frac{ln(S_t/K)}{K^2} P(t,\tau,K)dK\right) \tag{19}$$

Gao, Gao, & Song (2018) show that their RIX incorporates all higher-order cumulants and thus provides a natural downside tail-risk measure. Finally, we also include the classic VIX in our analysis:

$$VIX = \frac{2e^{r\tau}}{\tau}\left(\int_0^{S_t} \frac{1}{K^2} P(t,\tau,K)dK + \int_{S_t}^{\infty} \frac{1}{K^2} C(t,\tau,K)dK\right) \tag{20}$$

2.3 Robust estimation of risk-neutral moments

To obtain an accurate estimate of risk-neutral moments densely-spaced strike prices over a wide domain are needed. This requirement poses two problems: First, interpolation of the existing data, and second, the extrapolation outside the observed domain of option prices. The accuracy of the estimation of risk-neutral moments is thus dependent on the amount of available option data.

2.3.1 Interpolating the volatility surface

While there are many different approaches to obtain a smooth implied volatility surface, we focus on two approaches that are common in the literature and are easy to implement as they are available in most common software packages. The first approach uses a cubic smoothing spline and is the most common in the literature (e.g. Carr & Wu, 2009; Rehman & Vilkov, 2012). A smoothing spline fits multiple polynomials which connect smoothly to each other at knots. Smoothing splines require two parameters: The degree of the polynomials and the smoothness of the spline. We follow the most common approach to use cubic polynomials of order three.

We set the smoothing parameter to allow for an average variation of 0.01 in implied volatilities. In most software packages the smoothing factor is the maximum sum of squared errors that is admissible and the number of knots will be increased until this condition is achieved.

Second, we follow Aït-Sahalia & Lo (1998) and Song & Xiu (2016) and interpolate the implied volatility surface with a non-parametric kernel regression. Aït-Sahalia & Lo (1998) smooth the implied volatility surface with the local-constant or Nadaraya-Watson kernel estimator. However, in small samples the local-constant estimator tends to be inaccurate at the boundaries of the domain (e.g., Li & Racine, 2004) and thus artificially flattens the volatility surface. This leads to flatter than implied surfaces and should thus bias estimates of risk-neutral moments. Song & Xiu (2016) propose to use a local-linear kernel-regression instead to avoid under- or over-estimation of the implied volatility at the boundary. We use the local-linear estimator proposed by Stone (1977) and Cleveland (1979). The local-linear estimator has superior properties over the local-constant estimator (e.g., Fan, 1993). We choose to use a Gaussian kernel, but the choice of the kernel has little influence on the results (Aït-Sahalia & Lo, 1998). In contrast, the choice of the bandwidth is crucial to avoid over- or under-smoothing the implied volatility surface. We use leave-one-out cross-validation to select the bandwidth which is readily available in most software packages.

2.3.2 Extrapolating the volatility surface

Extrapolating the implied volatility surface is more challenging than interpolating it. The strikes of most traded options are located around the current stock price. In contrast, far out-of-the-money options are traded less frequently. Thus, we are observing the center portion of the risk-neutral distribution and under real-world data researchers will have make a choice about the extrapolation of the implied volatility surface. The extrapolation of the implied volatility surface is comparable

to adding tails to the risk-neutral probability density function. Extrapolating the implied volatility surface is a crucial step in the estimation of risk-neutral moments as higher order moments place a large weight on the tails of the risk-neutral distribution.

We analyze three different extrapolation techniques: First, we analyze risk-neutral moments if the volatility surface is not extrapolated similar to Conrad et al. (2013). If the volatility surface is not extrapolated researchers avoid making a decision about the shape of the tails of the risk-neutral distribution. However, under sparse option data the risk-neutral moment estimates will be biased as they focus only on the center of the distribution. For example, risk-neutral central skewness estimated from raw option data will estimate the skewness mostly from the center and shoulders of the risk-neutral distribution thus effectively creating a different measure of skewness. This limitation is particularly unfavorable if the risk-neutral moments of different assets should be compared, i.e. when sorting stocks on an estimate of risk-neutral skewness. Without extrapolation, stocks with more liquid option markets will generate skewness estimates which are mainly driven by the tail of the distribution while stocks with less liquid option markets generate skewness estimates which are constrained to the center of the risk-neutral distribution. Second, we extrapolate the implied volatility surface horizontally outside the known domain of strike prices, which is the most common procedure in the literature (e.g. Carr & Wu, 2009; Rehman & Vilkov, 2012). Assuming that the volatility surface is flat beyond the last observed strike is equivalent to assuming that the tails are normal. This assumption is unlikely to be true because the implied volatility surface typically exhibits a volatility skew and thus risk-neutral distributions are not normal. Moreover, horizontal extrapolation causes additional problems because it induces a kink in the volatility surface at the lowest and highest observed strike if the volatility surface is not flat at this point. The kink in the volatility surface is a discontinuity which can cause the risk-neutral PDF to become negative. Therefore, it is necessary to run an isotonic regression on the estimated cumulative density function to ensure that it is strictly increasing. Nevertheless, ex-

trapolating the implied volatility surface horizontally is a simple procedure which avoids any erratic behavior in the tails of the distribution making it a robust choice. Third, we use linear extrapolation, which extends the implied volatility surface as a linear function beyond the last observed options. G. J. Jiang & Tian (2007) show that extrapolating linearly is superior to the common method of flat extrapolation. If the volatility surface is upward sloping at the bounds the resulting tail will be heavier than the tail generated by a normal distribution. Linear extrapolation also has the advantage that the volatility surface is not kinked at the bounds helping to avoid a negative risk-neutral PDF. The slope at the bounds of the observed volatility surface is calculated with respect to the strike price.

2.3.3 Further considerations: Smoothing factor, degree of splines, and choice of variable

The estimation of risk-neutral moments is dependent on a variety of additional parameters. We opted to restrict this paper to the most relevant parameters, but for completeness we discuss these additional estimation options. To begin with, all our estimations are based on smoothing the implied volatility over strike prices. However, Shimko (1993) or Jurek (2014) propose to smooth the implied volatility surface over option delta calculated with the at-the-money implied volatility. Other authors (e.g. Carr & Wu, 2009; Stilger et al., 2016) choose to smooth the volatility surface over log-moneyness. In unreported results we also tested smoothing over deltas and over log-moneyness. In general, smoothing over strike prices and log-moneyness is less affected by narrow domains and noisy option data and the differences between both methods are neglectable. Smoothing over deltas has the theoretical advantage that it allows for most variation in close to-the-money implied volatilities where option prices are the most accurate. However, higher order moments are sensitive to the prices of far out-of-the-money options and changes in options close-to-the-money have relatively little influence on the risk-neutral moments. Therefore, the relatively

lower variation in the tail is a disadvantage leading to a slightly increased bias. Nevertheless, differences are small compared to total errors induced by narrow strike price domains and micro-structural noise. Furthermore, we vary the smoothing factor. In unreported results we tested smoothing factors allowing for an average approximation error in implied volatilities of 0.1 and 0.001, but found only minor differences in results. Finally, splines require that the degree of the polynomial functions between knots is defined. We repeated our simulation study with quartic splines and found little differences. In addition, we tested quartic smoothing splines with a single knot placed at-the-money as suggested by Birru & Figlewski (2012) but found no significant differences to cubic splines.

2.4 Results of the simulation study

To test the efficiency of the different extraction methods we extend the idea of Dennis & Mayhew (2009) and simulate option data. Simulated option data has the advantage that the true underlying distribution is known and estimates of moments can be compared to their true values. We test the efficiency of the different extraction methods by limiting the strike price domain of option prices, increasing the spacing between option prices, and by adding noise to the data.

2.4.1 Simulating option prices

Option prices generated by a stochastic-volatility-jump model similar to Bates (1996), where stock prices have stochastic volatility and a Poisson jump process with normally distributed jumps. An SVJ model gives us the flexibility to simulate non-normal distributions that are more realistic than the normal distribution assumed in the Black-Scholes model. The stochastic-volatility jump process is defined as:

$$\frac{dS_t}{S_t} = (r - \lambda\mu_J)dt + \sqrt{\sigma_t}dZ_t^1 + J_t dN_t \tag{21}$$

$$d\sigma_t = \kappa(\sigma_L - \sigma_t)dt + \nu\sqrt{\sigma_t}dZ_t^2 \tag{22}$$

$$corr(dZ^1, dZ^2) = \rho dt \tag{23}$$

$$prob(dN = 1) = \lambda dt \tag{24}$$

$$Z^1, Z^2 \sim N(0, 1) \tag{25}$$

$$ln(1 + J) \sim N(ln(1 + \mu_J) - 0.5\sigma_J^2, \sigma_J^2) \tag{26}$$

Where S is the stock price, r is the risk-free rate, σ is the volatility, κ is the mean reversion speed, σ_L is the long-run variance, ν is the volatility of volatility, Z^1 and Z^2 are two correlated standard normally distributed variables with correlation ρ, N is a Poisson distributed random variable with intensity λ where jumps, J_t, are log-normally distributed. The equations are discretized with an Euler scheme and option prices are based on 100'000 paths simulated with antithetic variables and moment matching. The volatility process is initialized with the long-run variance.

We generate two different scenarios: First, a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). And second, a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). In both scenarios we simulate option data with 90 days to maturity, time steps of half a day, an interest rate of 5%, and a stock price of 100 dollars. Based on the process we create option data that spans the domain from 1 dollar to 199 dollars in 0.5 dollar intervals. Thus, the domain half-width is 99% of the stock price and the spacing is 0.5 dollars or 0.5% of the stock price.[5]

[5]The exact parameters for replication in the standard scenario are: $\kappa = 2$, $\sigma_L = 0.05$, $\nu = 0.1$, $\rho = -0.6$, $\lambda = \mu_J = \sigma_J = 0$. The parameters are chosen to match the moments on a regular day, e.g. July 30 1997. Parameters for replication in the crisis scenario are: $\kappa = 0.5$, $\sigma_L = 0.3$, $\nu = 0.4$, $\rho = -0.95$, $\lambda = 1$, $\mu_J = -0.15$, $\sigma_J = 0.05$. The parameters are chosen to match the moments on a crisis day, e.g. October 12 2008 where the S&P500 had a return of -5.20%.

In each of the scenarios, the impact of narrow domains and micro-structural noise on the risk-neutral moments is tested following the methodology of Dennis & Mayhew (2009) and Bliss & Panigirtzoglou (2002). We also test the impact of wide strike price spacings. However, increased spacing between strike prices is relatively benign as filling the gaps is an interpolation exercise that is inherently easier then extrapolation. To keep the paper concise, we opt to present the results of restricted strike price domains and micro-structural noise in more depth. To test the effect of reduced domain width on risk-neutral moment estimates, the spacing is held constant at 50 cent and the dataset is truncated to a domain half-width of 10% of the stock price, from 90 to 110 dollars. Risk-neutral moments are then estimated from the truncated dataset. The domain half-width is then increased in 1% steps to a maximum of 99%. In each step, the risk-neutral moments are estimated from the truncated dataset. A similar procedure is also applied to test the effect of different strike price spacings. The domain half-width is held constant at 99% of the stock price and the spacing between strike prices is increased from 1 dollar (1%) to 10 dollars (10%) in 50 cent steps. The bias induced by market micro-structural noise is analyzed by perturbing the option prices in the dataset. Micro-structure noise is simulated by perturbing option prices as follows:

$$\widetilde{P}_i = P_i(1 + \theta\eta) \quad \eta \sim N(0,1) \tag{27}$$

where θ is increased from 1% to 10%. Micro-structural noise can be interpreted as the uncertainty about the true value of an option within the bid-ask spread. Assuming that perturbed option prices fall into the bid-ask bounds in 95% of all cases, the simulation can also be interpreted as varying the bid-ask spread from approximately 3.9% to 39%. After perturbing option prices, implied volatilities are again extracted and quantile and central moments are estimated and compared to the true moments. This procedure is repeated 1000 times for each noise level. To test the influence of micro-structural noise under realistic conditions we restrict the domain from 80 to

120 dollar and keep the spacing at 2.5 dollar.

In each step we interpolate and extrapolate the implied volatility surface using three different methods: First, we use a cubic smoothing spline with horizontal extrapolation (spline-flat), second, a cubic smoothing spline with linear extrapolation (spline-linear), and third, a local-linear kernel regression with linear extrapolation (kernel-linear). The smoothing factor of the splines allows for an average error in implied volatilities of 0.01 and the bandwidth of the local-linear kernel regression is determined through leave-one-out cross-validation. Our results are benchmarked against a plain-vanilla implementation of Bakshi et al. (2003) with no inter- or extrapolation as used by Dennis & Mayhew (2009) and Conrad et al. (2013). In addition to the central risk-neutral moments of Bakshi et al. (2003), we also estimate quantile moments based on the model-free approach and the naïve Black-Scholes approximation. Quantile volatility is defined as the inter-quartile range between the 75^{th} and 25^{th} percentile. Quantile skewness is measured at the 10^{th} and 90^{th} percentile ($p = 0.9$) and quantile kurtosis is estimated with an outer range between the 95^{th} and 5^{th} percentile and an inner range between the 75^{th} and 25^{th} percentile ($p = 0.95$ and $q = 0.75$). Estimation errors are calculated in percent of the true values because the estimates of quantile and central moments are not on the same scale.

2.4.2 Results of risk-neutral moments

We begin by analyzing the bias in risk-neutral moments introduced by an insufficiently wide domains. The results are presented in Table 11 and Figure 2. Panels A & B of Figure 2 show the efficiency of different inter- and extrapolation techniques. In the calm scenario all estimation methods provide accurate estimates of central volatility. However, the plain-vanilla implementation (BKM-raw) converges more slowly to the true value in comparison to the other methods. The estimation errors are higher for estimates of central skewness. Notably, the two methods that rely on

linear extrapolation deliver a speedy convergence to the true value of skewness. In contrast, a smoothing spline approach with horizontal extrapolation estimates skewness initially with an absolute error of 40% and converges at a domain half-width of 60%. Moreover, no inter- or extrapolating yields an error of 60% initially. The results are similar for kurtosis. Linear extrapolation delivers lower initial errors and faster convergence to the true value. This finding becomes even more apparent in the crisis scenario (Panel B). While both extrapolation methods only converge to the true value of central moments at a domain half-width of approximately 90%, linear extrapolation has a significantly lower estimation bias if the domain of option prices spans only a narrow domain. This observation is in line with the results of G. J. Jiang & Tian (2007) and can be attributed to the different tail shapes of the risk-neutral density that are caused by the different extrapolation methods. Horizontal extrapolation assumes a flat volatility surface beyond the last available option prices and thus essentially adds the tail of a normal distribution. In contrast, linear extrapolation leads to heavier than normal tails if the surface is upward sloping at the bounds, leading to a better approximation of the underlying distribution. Moreover, the choice of the smoothing method, cubic smoothing splines or a kernel regression, does not have a large influence on the bias of the estimation caused by an insufficiently wide domain. Panels C & D show a comparison of quantile and central moments based on a kernel regression. Notably, the model-free implementation of quantile moments delivers accurate estimates even for narrow domains under both scenarios for all moments. This finding becomes especially apparent when the performance of the model-free quantile moments estimator is compared to the common Bakshi et al. (2003) cubic spline implementation with horizontal extrapolation. This finding is not coincidental and is related to the construction of quantile and central moments. Central moments require to find $E[X]$, $E[X^2]$, $E[X^3]$, etc. The higher the order of the desired moment, the larger is the weight on the tail of the distribution. Thus, the estimation of higher order central moments is almost certainly biased until the entire PDF is recovered. In contrast, quantile moments only require the correct estimation of the

relevant quantiles and thus can deliver accurate estimates even when option prices do not reveal the entire PDF. Moreover, the naïve Black-Scholes approximation does not converge to the true value of any quantile moment; however, it nearly instantly converges to a stable value. In addition to the results in Figure 2, Table 11 shows the results in tabular from for the domain half-widths of 10%, 50%, and 80%. In addition, Table 11 also contains estimates of quantile moments based on cubic splines with flat extrapolation. Even under a horizontal extrapolation quantile moments have a lower bias under narrow domains than their central counterparts. However, the estimation of model-free quantile skewness from a spline-linear model does not converge to the true value of quantile skewness in the standard scenario. This finding is most likely due to the kink in the implied volatility surface caused by the horizontal extrapolation leading to numerical problems when differentiating the option-price surface to obtain the CDF. If the relevant quantile happens to be close to the kink, estimates of quantile moments become unstable and biased. We therefore recommend to estimate quantile moments from linearly extrapolated volatility surfaces. Moreover, Table 11 also contains estimates obtained from a local-constant kernel regression comparable in magnitude to the bias of the cubic smoothing spline with horizontal extrapolation.

Table 11: Errors from truncated domain half-width - SVJ

This table shows the approximation errors of option-implied moments for different strike price domain half-widths in percent. The strike price spacing is held constant at 50 cents. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). Panel B displays the results of a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). For each of the scenarios risk-neutral moments are estimated under different domain half-widths (10%, 50%, & 80%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Domain half-width	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	12.35	0.06	0.08	61.45	1.58	0.81	68.22	2.54	0.46
BKM-spline-flat	2.56	0.08	0.00	38.34	2.55	1.63	30.45	1.36	0.31
BKM-spline-linear	0.22	0.08	0.00	3.17	2.40	1.25	2.51	0.94	1.05
BKM-kernel-linear	0.20	0.00	0.01	2.01	0.92	0.45	2.90	0.91	0.57
BKM-lckernel-linear	2.66	0.01	0.01	37.26	0.35	0.36	29.40	0.08	0.10
Quantile Moments									
MFree-kernel-linear	0.11	0.11	0.11	6.00	0.61	0.64	0.56	0.09	0.09
MFree-spline-flat	0.55	0.17	0.33	17.96	16.16	12.94	0.24	0.57	0.06
MFree-lckernel-linear	0.25	0.32	0.32	20.70	1.71	1.71	0.34	0.39	0.39
BS-kernel-linear	7.53	7.53	7.53	11.35	13.02	13.02	1.99	1.09	1.09
BS-spline-flat	7.55	7.37	7.49	35.06	15.54	14.72	6.73	0.65	0.98
BS-lckernel-linear	7.89	7.89	7.89	35.49	13.23	13.23	7.41	1.72	1.72
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	44.66	8.37	0.72	94.86	35.40	5.84	94.50	58.44	14.65
BKM-spline-flat	19.30	4.37	0.40	71.28	21.02	3.48	70.66	39.71	8.92
BKM-spline-linear	1.92	0.97	0.19	1.41	6.22	1.78	8.79	14.07	4.20
BKM-kernel-linear	0.77	1.00	0.19	6.89	6.36	1.87	15.59	14.30	4.41
BKM-lckernel-linear	19.21	4.44	0.39	70.65	21.41	3.37	71.06	40.80	9.28
Quantile Moments									
MFree-kernel-linear	0.70	0.04	0.04	2.63	0.32	0.36	2.79	0.20	0.19
MFree-spline-flat	14.84	0.05	0.75	17.77	0.24	0.64	22.97	11.26	0.36
MFree-lckernel-linear	14.53	0.13	0.21	18.30	0.23	0.23	22.55	11.32	0.12
BS-kernel-linear	28.35	28.01	28.01	18.40	18.00	18.00	10.61	10.76	10.55
BS-spline-flat	14.84	27.88	27.67	73.32	18.21	18.22	22.97	14.80	10.22
BS-lckernel-linear	14.27	27.21	27.21	72.63	18.01	18.01	22.47	14.41	10.03

Figure 2: Errors in the approximation of implied moments by domain half-width

This figure shows the approximation errors of option-implied moments for different strike price domain half-widths. The strike price spacing is held constant at 50 cents. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: Panels A & C show the results of a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). Panels B & D display the results of a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). For each of the scenarios risk-neutral moments are estimated under different domain half-widths. Panel A & B show the performance of different inter- and extrapolation techniques. Risk-neutral moments in these panels are based on Bakshi et al. (2003, BKM). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation or a non-parametric local-linear kernel regression similar to Song & Xiu (2016). Panel C & D compare the standard approach of Bakshi et al. (2003) with quantile-moments estimated either model-free (QuantMF) or from a naïve Black-Scholes approximation (QuantBS). The error in the plots is truncated at 80%.

Panel A: Standard Scenario - Different Inter- and Extrapolation

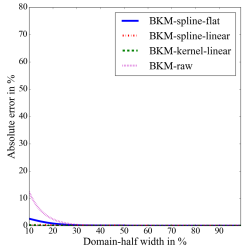

(a) Central volatility

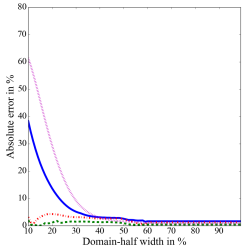

(b) Central skewness

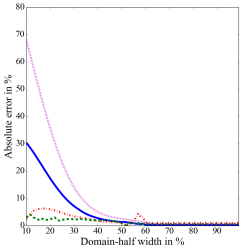

(c) Central kurtosis

Panel B: Crisis Scenario - Different Inter- and Extrapolation

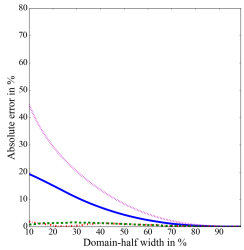

(d) Central volatility

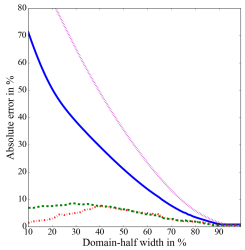

(e) Central skewness

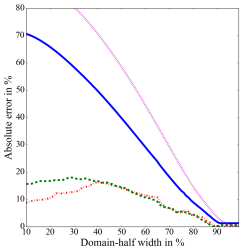

(f) Central kurtosis

continued on the next page

Figure 2 continued

Panel C: Standard Scenario - Different Estimators

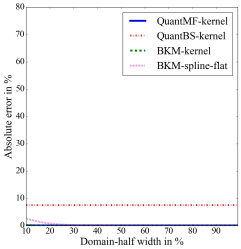

(g) Quantile and central volatility

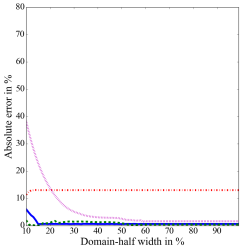

(h) Quantile and central skewness

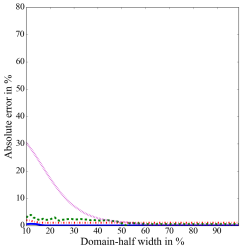

(i) Quantile and central kurtosis

Panel D: Crisis Scenario - Different Estimators

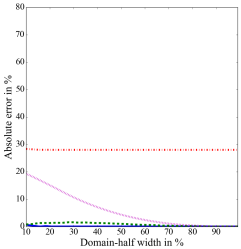

(j) Quantile and central volatility

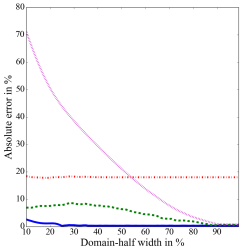

(k) Quantile and central skewness

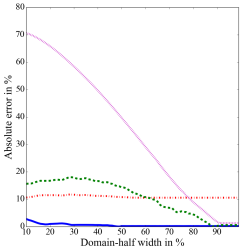

(l) Quantile and central kurtosis

The sensitivity of estimates to micro-structural noise is shown in Figure 3 and Table 12. Panels A and B show the standard deviation of estimates from different inter- and extrapolation methods to micro-structural noise under a standard (Panel A) and crisis scenario (Panel B). We can observe that a plain-vanilla implementation of Bakshi et al. (2003), without inter- and extrapolation, is the most robust to noise for all moments and for both scenarios. Moreover, a cubic spline with linear extrapolation is very sensitive to micro-structural noise. For example under the spline-linear method, perturbing option prices with a standard deviation of 5% leads to a standard deviation of central skewness estimates of 147.09% in the standard scenario (Panel A of Table 12). In contrast, the commonly used method of cubic smoothing splines with horizontal extrapolation fares relatively well for all moments in all scenarios. For example,

even under micro-structural noise with a standard deviation of 10%, the standard deviation of central skewness estimates is just 6.60% under the crisis scenario (Panel B of Table 12). Estimates of the kernel regression technique are reasonably robust to micro-structural noise, but are slightly more affected than estimates using cubic smoothing splines with horizontal extrapolation. Moreover, the local-constant kernel regressions with linear extrapolation deliver relatively stable estimates and perform even slightly better than the cubic smoothing splines with horizontal extrapolation in some cases. The standard deviations of quantile moments are comparable to their central counterparts. Under the standard scenario (Panel C in Figure 3) the quantile moments estimated from a naïve Black-Scholes approximation are the least affected by noise. In contrast, model-free quantile moments tend to pick up more micro-structural noise. In comparison with central moments, naïve Black-Scholes quantile moments are slightly less affected by noise, whereas model-free quantile moments are more affected by micro-structural noise. Under the crisis scenario, the differences between estimation methods are smaller. Central moments estimated with smoothing splines and horizontal extrapolation are the most robust to micro-structural noise. However, the differences between moments are relatively small.

The previous results reveal a bias-variance trade-off.[6] Methods with a relatively low bias under narrow domains, i.e. central moments estimated from a smoothed volatility surface with linear extrapolation, are relatively strongly affected by micro-structural noise leading to a high variance in the estimates. In contrast, methods with a high bias under narrow domains, i.e. a plain-vanilla implementation of Bakshi et al. (2003), have a low variance of estimates even under high levels of micro-structural noise. The most common method in the literature, central moments estimated from cubic smoothing splines with flat extrapolation, has a relatively high bias under narrow domain-width, but is relatively little affected by micro-structural noise. These properties make this method attractive for datasets with relatively wide domains,

[6] The bias-variance trade-off exists under real-world data. Asymptotically all methods are unbiased.

such as single stocks with highly liquid option markets or S&P 500 index options. In contrast, central moments estimated from datasets with narrow domains, such as single stocks with relatively illiquid options, likely carry a high bias and under or over estimate risk-neutral moments. At the same time, the relatively high level of micro-structural noise in prices of illiquid single stock options will have only a minor affect on the estimates of risk-neutral moments based on the spline-flat method. An opposite result is delivered by quantile moments estimated from a linearly extrapolated volatility surface. While the estimates carry a low bias, they are affected by micro-structural noise to a greater extent. Quantile moments estimated from a linearly extrapolated volatility surface have a significantly lower bias compared to the standard method, especially under narrow domains.

Table 12: Errors from micro-structural noise - SVJ

This table shows the standard deviation of estimates of option-implied moments for different levels of micro-structural noise in percent. The strike price spacing is set to 2.5 dollar and strikes cover a range from 80% to 120% of the strike price. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). Panel B displays the results of a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). Micro-structure noise is simulated by perturbing option-prices by a percent of their value. Percentages are drawn randomly from a normal distribution with a standard deviation of either 1%, 5%, or 10%. For each level of micro-structural noise risk-neutral moments are estimated 1000 times. Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Noise	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.14	0.72	1.42	0.50	2.47	5.04	0.34	1.66	3.28
BKM-spline-flat	0.15	1.90	18.13	0.57	9.92	16.86	0.51	2.72	7.75
BKM-spline-linear	0.20	22.47	104.44	2.43	147.09	339.19	3.31	312.15	565.41
BKM-kernel-linear	0.20	0.80	1.58	2.11	7.25	18.54	2.84	10.03	34.48
BKM-lckernel-linear	0.15	0.75	1.49	0.64	3.80	6.86	0.53	2.87	5.27
Quantile Moments									
MFree-kernel-linear	2.38	6.14	9.15	10.86	22.05	28.64	2.55	7.62	15.73
MFree-spline-flat	0.99	13.27	22.13	4.79	54.83	90.14	1.15	17.32	696.33
MFree-lckernel-linear	5.73	8.57	12.54	25.27	32.45	43.07	5.90	10.16	31.29
BS-kernel-linear	0.31	1.05	1.89	0.80	2.69	5.24	0.34	1.25	2.50
BS-spline-flat	0.21	2.22	4.93	0.65	4.96	25.13	0.26	2.13	11.82
BS-lckernel-linear	0.37	1.28	2.46	0.95	3.56	6.56	0.37	1.44	2.62
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.09	0.46	0.92	0.13	0.67	1.38	0.05	0.25	0.49
BKM-spline-flat	1.00	5.73	6.59	0.55	7.11	6.60	0.75	3.22	3.98
BKM-spline-linear	20.36	80.89	67.27	34.09	1.02e3	1.12e3	44.56	1.24e3	1.77e4
BKM-kernel-linear	2.50	9.17	19.22	6.05	15.05	24.77	8.15	19.02	25.25
BKM-lckernel-linear	0.24	1.06	2.03	0.46	2.22	4.42	0.27	1.36	2.57
Quantile Moments									
MFree-kernel-linear	3.06	8.62	12.92	2.90	9.18	14.35	3.67	14.69	28.66
MFree-spline-flat	8.25	18.71	21.30	14.11	30.86	31.27	8.13	67.96	103.64
MFree-lckernel-linear	5.86	16.52	21.45	6.50	12.83	15.97	5.69	16.44	23.56
BS-kernel-linear	0.95	5.01	30.41	3.00	8.72	14.47	2.92	7.96	11.33
BS-spline-flat	0.80	3.12	7.00	1.60	4.18	5.84	0.52	1.62	32.55
BS-lckernel-linear	0.50	2.19	3.81	0.60	3.00	5.70	0.13	0.96	1.25

Figure 3: Errors in the approximation of implied moments for different levels of micro-structural noise

This figure shows the standard deviation of estimates of option-implied moments for different levels of micro-structural noise. The strike price spacing set to 2.5 dollar and strikes cover a range from 80% to 120% of the strike price. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: Panels A & C show the results of a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). Panels B & D display the results of a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). Micro-structure noise is simulated by perturbing option-prices by a percent of their value. Percentages are drawn randomly from a normal distribution with a standard deviation ranging from 1% to 10%. For each level of micro-structural noise risk-neutral moments are estimated 1000 times. Panel A & B show the performance of different inter- and extrapolation techniques. Risk-neutral moments in these panels are based on Bakshi et al. (2003, BKM). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation or a non-parametric local-linear kernel regression similar to Song & Xiu (2016). Panel C & D compare the standard approach of Bakshi et al. (2003) with quantile-moments estimated either model-free (QuantMF) or from a naïve Black-Scholes approximation (QuantBS). The standard deviation of the estimates in the plots is truncated at 100%.

Panel A: Standard Scenario - Different Inter- and Extrapolation

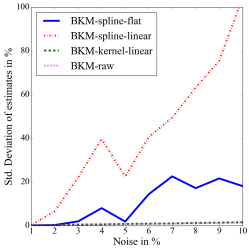

(a) Central volatility

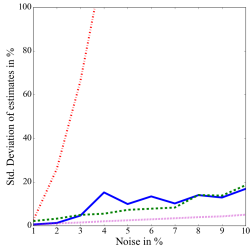

(b) Central skewness

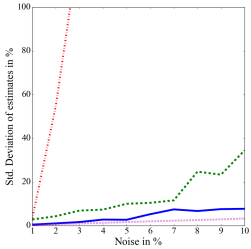

(c) Central kurtosis

Panel B: Crisis Scenario - Different Inter- and Extrapolation

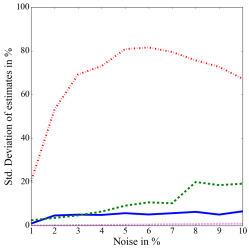

(d) Central volatility

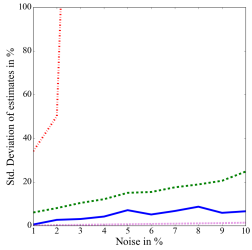

(e) Central skewness

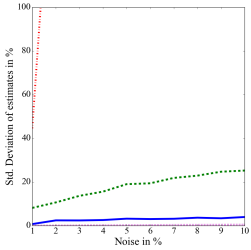

(f) Central kurtosis

continued on the next page

Figure 3 continued

Panel C: Standard Scenario - Different Estimators

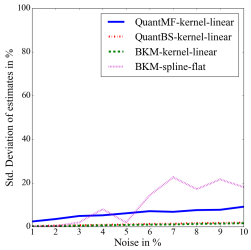

(g) Quantile and central volatility

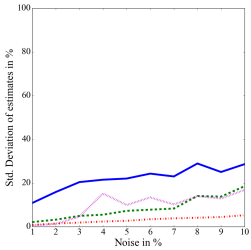

(h) Quantile and central skewness

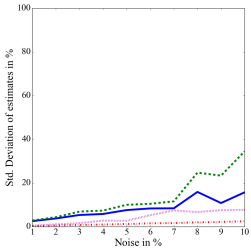

(i) Quantile and central kurtosis

Panel D: Crisis Scenario - Different inter- and Extrapolation

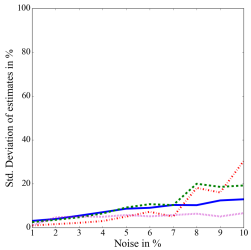

(j) Quantile and central volatility

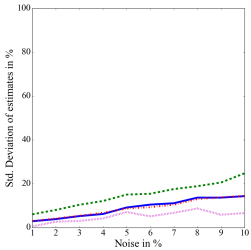

(k) Quantile and central skewness

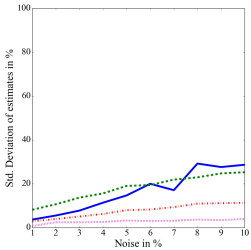

(l) Quantile and central kurtosis

The approximation errors for different domain spacings (Table 13) are generally smaller than the approximation errors caused by a small domain. The error is also easier to control as smoothing over strike prices is an interpolation task whereas extending the domain requires extrapolation. The approximation errors of all methods do not seem to be systematically biased by strike price spacings and are negligible in most scenarios. The only exception are the estimation errors of the plain-vanilla implementation of Bakshi et al. (2003) without inter- or extrapolation, which increase slightly with larger strike price spacings.

2.4.3 Results of RIX, VIX, and SVIX

In addition to the classic risk-neutral moments and their quantile equivalents we also analyze the sensitivities of the RIX (Gao, Gao, & Song, 2018; Gao, Lu, & Song,

Table 13: Errors from different strike price spacing - SVJ

This table shows the approximation error of option-implied moments for different strike price spacings in percent. The domain-half width is set to 99% of the current stock price. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with central (quantile) volatility of 0.23 (0.14), skewness of -0.89 (-0.20), and kurtosis of 4.72 (2.61). Panel B displays the results of a crisis scenario with central (quantile) volatility of 0.64 (0.30), skewness of -2.27 (-0.55), and kurtosis of 10.39 (3.07). For each of the scenarios risk-neutral moments are estimated under different strike price spacings (1%, 2%, & 5%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Spacing	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.08	0.13	0.44	0.81	0.34	10.53	0.46	0.38	2.93
BKM-spline-flat	0.00	0.01	0.02	1.63	1.01	1.51	0.31	0.49	0.56
BKM-spline-linear	0.00	0.01	0.04	1.25	0.02	0.84	1.05	3.32	1.85
BKM-kernel-linear	0.01	0.02	0.09	0.45	0.47	0.83	0.57	0.70	0.62
BKM-lckernel-linear	0.00	0.00	0.02	0.34	0.35	0.35	0.08	0.11	0.20
Quantile Moments									
MFree-kernel-linear	0.11	0.55	0.25	0.64	0.76	4.03	0.09	0.58	0.27
MFree-spline-flat	0.33	0.28	0.26	12.94	6.78	12.09	0.06	0.71	0.07
MFree-lckernel-linear	0.02	0.51	7.48	15.82	3.87	6.48	0.47	0.67	12.66
BS-kernel-linear	7.53	7.54	7.63	13.02	13.12	13.77	1.09	1.11	1.15
BS-spline-flat	7.49	7.52	7.53	14.72	13.62	14.24	0.98	0.94	0.96
BS-lckernel-linear	7.45	7.91	10.22	14.11	14.62	17.33	1.58	1.31	4.60
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.04	0.07	0.04	0.59	0.51	0.19	0.33	0.27	0.36
BKM-spline-flat	0.08	0.07	0.09	0.82	0.85	1.14	1.29	1.36	2.30
BKM-spline-linear	0.05	0.05	0.03	0.47	0.52	0.47	0.01	0.14	0.05
BKM-kernel-linear	0.06	0.05	0.01	0.62	0.59	0.46	0.56	0.45	0.11
BKM-lckernel-linear	0.05	0.05	0.03	0.64	0.60	0.62	1.67	1.60	1.94
Quantile Moments									
MFree-kernel-linear	0.04	0.06	0.16	0.33	0.50	0.83	0.20	0.04	0.18
MFree-spline-flat	0.25	0.45	1.04	0.92	0.88	0.14	0.14	0.31	0.97
MFree-lckernel-linear	0.76	0.37	3.90	1.40	5.74	5.79	0.37	2.19	1.33
BS-kernel-linear	28.01	28.01	27.98	18.00	18.01	18.13	10.55	10.54	10.48
BS-spline-flat	27.95	28.00	27.77	18.20	18.19	18.25	10.38	10.42	10.22
BS-lckernel-linear	27.17	27.17	29.24	18.13	17.93	20.28	9.51	9.29	11.41

2018), the classic VIX, and the SVIX (Martin, 2017) to changes in the domain width, micro-structural noise, and strike price spacing with the same methodology. The results are presented in Table 14.

Panel A shows the effect of reduced domain half-widths on the estimates of the VIX, RIX, and SVIX. The results of the VIX and SVIX are comparable to the option-implied volatility measure of Bakshi et al. (2003). All three measures provide accurate estimates at a domain half-width of 50% under both the standard and crisis scenario. Notably, the RIX requires a much narrower domain than corresponding central moments, e.g. skewness or kurtosis, to deliver accurate an unbiased estimate. For example, RIX estimated with a local-linear kernel regression at a domain half-width of 50% has an error of 0.01% under the crisis scenario. In contrast, the risk-neutral kurtosis measure of Bakshi et al. (2003) is estimated with an error of 14.30%. The RIX should provide a complete measure of all higher order moments and is therefore a suitable candidate for studies that aim to estimate downside risk. The differences between estimation methodologies mirrors the previous findings. Linear extrapolation leads to a lower bias under narrow domain half-width, but the differences are less severe as the VIX, RIX, and SVIX converge faster to their true values.

Panel B shows the impact of micro structural noise on the VIX, RIX, and SVIX estimates. The results are approximately on par with the previous findings. Linear extrapolation combined with a local-linear kernel regression leads to a higher sensitivity towards micro-structural noise, especially in the crisis scenario. In contrast, the spline-flat method is less affected by micro structural noise under the the crisis scenario, but slightly more affected under the standard scenario. The best performing methodology is the local-constant kernel regression of Aït-Sahalia & Lo (1998) with linear extrapolation. For example, under the crisis scenario, the percentage error of the RIX estimate from a local-constant kernel regression has a standard deviation of 32.49% whereas the standard deviation from the spline-flat method is 37.84%.

Panel C displays the effect of changes in strike price spacing on the estimates of VIX, RIX, and SVIX. Changes in strike price spacing are negligible compared to other errors. The error induced by a 5% strike price spacing on is less than 1% for all methodologies for VIX, RIX, and SVIX if the volatility surface is interpolated. Even if the volatility surface is not interpolated the maximum error is only 5.13% for the RIX estimated under the standard scenario. In conclusion, the local-constant kernel regression estimator with linear extrapolation appears to offer the best trade-off between bias and variance. Local-linear kernel regression leads to reduced bias but increased variance. Cubic smoothing splines with horizontal extrapolation lead to a similar bias as the local-constant kernel regression with linear extrapolation but their estimates are more affected by micro-structural noise.

2.5 Empirical estimation

We aim to verify the efficiency of the extraction methods also empirically by comparing summary statistics for different inter- and extrapolation techniques as well as for quantile and central moments.

2.5.1 Data and estimation

The analysis focuses on S&P 500 options. Daily option data is sourced from the OptionMetrics price database (provided by Wharton Research Data Service) for the sample time period from January 1996 to December 2017.

OptionMetrics provides two datasets, raw option data and a pre-smoothed volatility surface. For most applications the use of the volatility surface will lead to stable results and computations will be less expensive. However, to provide a true comparison between the previous methods, we estimate option-implied moments from raw

Table 14: Results of VIX, RIX, and SVIX

This table shows the approximation errors of the VIX, RIX, and SVIX measures for different strike price domain half-widths (Panel A), micro-structural noise (Panel B), and different strike price spacings (Panel C) in percent. The return process is based on a stochastic volatility jump process (Bates, 1996). Option-implied moments are estimated under two different scenarios: A standard scenario and a crisis scenario. For each of the scenarios risk-neutral moments are estimated under different domain half-widths levels, levels of micro-structural noise, and strike price spacings. Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Panel A and C display the absolute percentage error compared to the true estimate. Panel B shows the standard deviation of the percentage error. All results are shown in percent.

	VIX			**RIX**			**SVIX**		
Panel A: Domain half-width									
	10%	**50%**	**80%**	**10%**	**50%**	**80%**	**10%**	**50%**	**80%**
Standard scenario									
kernel-linear	0.41	0.00	0.00	0.34	0.01	0.01	0.67	0.00	0.00
lckernel-linear	4.01	0.00	0.00	25.91	0.00	0.00	1.96	0.00	0.00
spline-flat	4.05	0.02	0.02	26.11	0.30	0.24	1.98	0.04	0.03
spline-linear	1.06	0.02	0.02	3.76	0.30	0.23	1.03	0.04	0.03
raw	21.92	0.00	0.00	64.35	0.04	0.00	19.50	0.00	0.00
Crisis scenario									
kernel-linear	0.41	0.00	0.00	0.34	0.01	0.01	0.67	0.00	0.00
lckernel-linear	4.01	0.00	0.00	25.91	0.00	0.00	1.96	0.00	0.00
spline-flat	4.05	0.02	0.02	26.11	0.30	0.24	1.98	0.04	0.03
spline-linear	1.06	0.02	0.02	3.76	0.30	0.23	1.03	0.04	0.03
raw	21.92	0.00	0.00	64.35	0.04	0.00	19.50	0.00	0.00

continued on the next page

Table 14 continued

	VIX			RIX			SVIX		
Panel B: Noise									
	1%	**5%**	**10%**	**1%**	**5%**	**10%**	**1%**	**5%**	**10%**
Standard scenario									
kernel-linear	0.36	1.58	3.12	1.92	6.60	31.81	0.34	1.53	2.83
lckernel-linear	0.31	1.55	3.05	0.48	2.92	5.35	0.32	1.56	3.07
spline-flat	0.31	1.58	3.46	0.46	2.62	79.31	0.31	1.59	3.42
spline-linear	0.39	365.81	533.75	2.55	2.06e4	2.83e4	0.33	43.88	153.16
raw	0.30	1.47	2.95	0.37	1.87	3.61	0.30	1.49	2.99
Crisis scenario									
kernel-linear	5.90	19.20	70.19	78.89	287.17	1.72e3	1.55	4.94	10.75
lckernel-linear	0.64	2.80	5.38	3.82	16.91	32.49	0.40	1.81	3.51
spline-flat	0.67	3.04	6.14	3.85	18.00	37.84	2.71	3.65	6.72
spline-linear	56.29	445.50	421.62	1.15e3	1.33e4	1.18e4	10.42	142.22	363.67
raw	0.26	1.28	2.54	0.39	1.96	3.90	0.26	1.27	2.51
Panel C: Strike price spacing									
	1%	**2%**	**5%**	**1%**	**2%**	**5%**	**1%**	**2%**	**5%**
Standard scenario									
kernel-linear	0.01	0.02	0.08	0.01	0.01	0.01	0.01	0.02	0.10
lckernel-linear	0.01	0.01	0.05	0.00	0.02	0.15	0.01	0.01	0.05
spline-flat	0.00	0.02	0.10	0.24	0.20	0.14	0.04	0.03	0.06
spline-linear	0.00	0.02	0.10	0.23	0.20	0.13	0.04	0.03	0.06
raw	0.14	0.27	1.22	0.11	0.22	5.13	0.14	0.28	1.91
Crisis scenario									
kernel-linear	0.00	0.01	0.01	0.10	0.08	0.06	0.01	0.00	0.02
lckernel-linear	0.01	0.01	0.01	0.24	0.21	0.12	0.01	0.00	0.00
spline-flat	0.03	0.03	0.07	0.24	0.24	0.43	0.01	0.01	0.04
spline-linear	0.01	0.02	0.04	0.09	0.10	0.15	0.01	0.01	0.04
raw	0.02	0.04	0.03	0.01	0.02	0.48	0.03	0.06	0.38

option data. The estimation of risk-neutral moments from raw option price data poses a number of additional challenges. To begin with, the goal of most studies is to compare option-implied information at a constant maturity. However, real-world option data is often not available at the target maturity. Especially, when working with single stock options this poses a challenge as options mature only once a month. We tackle this challenge by first interpolating implied volatility across log-moneyness and then interpolate linearly between volatilities. Moreover, the implied volatility of an option can be computed either at its ask, bid, or mid price. We use the OptionMetrics provided implied volatilities which are estimated at the mid price and account for dividends.

Another challenge is to ensure that sufficient option data is available. Many studies (e.g. Conrad et al., 2013) filter out options with zero volume or zero open-interest to avoid stale prices. In our experience, strong filters, e.g. requiring strictly positive volume, lead to more noisy estimates especially when combined with an extrapolation technique. This might seem counter-intuitive at first, as strong filters should remove stale prices. However, strong filters also change the available data from period to period. This effect is particularly severe if the surface is extrapolated. Removing a deep out-of-the-money option can have a drastic effect, especially on option-implied information that rely on the tails of the distribution such as skewness. We do not employ any volume or open-interest filters. Instead, we filter out all options with a delta absolutely smaller than 0.1 to avoid erratic behavior in the tails of the distribution. However, the choice of filter should always be tailored to the research application of the extracted option-implied information.

Moreover, we rely only on out-of-the-money options, thus all options with a strike below the current stock price are put options and all options with strikes above the current strike prices are calls. Therefore, the call and put volatilities have to be joined at-the-money. For the S&P 500 at-the-money put and calls typically have identical volatilities.[7]

[7] Note that this is non-trivial for single-stock options. Single-stock options are American options and thus calls and puts do not necessarily have the same volatility (see Cremers & Weinbaum,

Our estimation approach is as follows: We begin by filtering out all in-the-money options and options with a negative implied volatility as well as all options with an absolute delta smaller than 0.1. We restrict our sample to options between with a maturity 1 and 60 days and target a constant maturity of 30 days. Implied volatilities are first interpolated for each maturity across strikes the implied volatility surface over a strike price interval from -99% to $+99\%$ with a spacing of 1 cent between strikes. We create volatility surfaces using the two most promising techniques from section 2.4: A local-linear kernel regression with linear extrapolation and a cubic smoothing spline with horizontal extrapolation. From the volatility surface we obtain risk-neutral quantile moments using the previously described methodology and the moments of Bakshi et al. (2003).

Summary statistics of the S&P 500 options are provided in Table 15. The S&P 500 has a large and liquid option market with an average contract volume of 371.83 contracts traded daily per option. Moreover, on average the open interest is 4064.11 contracts per options and the average domain width of the S&P 500 is 103.85% and the spacing between strikes is on average 1%.

2.5.2 Comparison of quantile and central moments

Summary statistics for option-implied moments and the VIX, RIX, and SVIX, are provided in Table 16. Panel A shows that the risk-neutral distribution of the S&P 500 is on average left-skewed and leptokurtic. This observation matches the findings of other studies in the literature (e.g. Neumann & Skiadopoulos, 2013). However, we can observe notable differences between the different estimation techniques as well as between central and quantile moments. The summary statistics of quantile moments estimated from a kernel-linear and spline-flat method are similar. The spline-flat method leads to slightly (absolutely) higher estimates of quantile-skewness and quantile-kurtosis although the differences are not statistically significant. Fur-

2010). Naively joining put and call volatilities at-the-money will lead to a jump at-the-money.

Table 15: Summary statistics of options

This table displays descriptive statistics of the sample of S&P 500 options. The sample is restricted to options with an absolute delta larger or equal to 0.1. Days to expiration measures the days until the option expires. The domain width is calculated for each day-maturity combination by subtracting the smallest strike from the largest strike and dividing this range by the closing price of the S&P 500 (SPX). Spacing is the average spacing between strike prices for each day-maturity combination divided by the closing price of the S&P 500. Domain and spacing are displayed in percent. The sample spans the period between January 1996 and December 2017.

	Mean	*Std.*	*Med.*	*Min*	*Max*
Days to expiration	28.80	16.45	28.00	2.00	59.00
Strike Price	1725.46	517.96	1770.00	50.00	3500.00
Implied Volatility	0.33	0.32	0.24	0.02	3.00
Volume	371.83	1960.44	0.00	0.00	200777.00
Open interest	4064.11	13697.76	100.00	0.00	370769.00
Delta	0.15	0.57	-0.00	-1.00	1.00
Domain width in % of SPX	103.85	43.17	93.85	26.30	266.06
Spacing in % of SPX	1.07	0.47	1.07	0.38	2.62

thermore, the summary statistics for central volatility are similar for both estimation methods. In contrast, central skewness and central kurtosis are heavily affected by the choice of the estimation method. Central moments estimated with the kernel-linear method show that the implied distribution of the S&P 500 is more left skewed and leptokurtic than if central moments are estimated with the spline-flat method. For example, the spline-flat method estimates the mean option-implied skewness as -0.79. This result is consistent with estimates in the literature that use a similar technique, e.g. -0.91 for 60 day maturity options in the study of Neumann & Skiadopoulos (2013). In contrast, the kernel-linear method estimates central skewness to be -0.97. However, central skewness estimates from the kernel-linear method are more volatile (standard deviation of 0.49%) than those obtained from the spline-flat method (standard deviation 0.37%). Moreover, the results of quantile skewness and kurtosis results are more robust to the modeling choice of the volatility surface and deliver consistent results for both estimation methods. In contrast, central moments use probability weighting of outcomes and are very sensitive to small shifts in the tail probability mass. For example, central excess kurtosis has a mean of 1.04 and takes

a maximum value of 6.85 if it is estimated with the spline-flat method. In contrast, it has a mean of 1.93 and takes a maximum value of 59.55 if it is estimated with the kernel-linear method. The results of the VIX, and SVIX show that their distribution is relatively unaffected by the choice of estimation method. This finding is in line with the previous results from the simulation study.

The observed results mirror the bias-variance trade-off observed in the simulation study. Moments obtained from the spline-flat method likely carry a bias but are less volatile than moments obtained from the kernel-linear method. Panels B and C show the correlations between the time series of moments. The correlations between quantile and central moments estimated from the kernel-linear (spline-flat) method are 75.05% (71.40%), 60.10% (73.84%), and 38.37% (48.74%) for volatility, skewness, and kurtosis, respectively. The correlations and show that quantile and central moments capture similar properties of the risk-neutral distribution. Moreover, the correlation between central and quantile kurtosis is notably lower than for volatility and skewness in Panel C. Quantile kurtosis of Ruppert (1987) is a pure measure of tail-weight, while central kurtosis measures tail-weight and peakedness simultaneously. Furthermore, the correlation pattern within quantile and central moments matches in sign. The findings indicate that quantile and central moments measure similar attributes of the risk-neutral distribution. The correlation of the VIX, RIX, and SVIX also provides interesting insights. SVIX and quantile volatility are all highly correlated with a correlation coefficient of 94.54% (94.24%) estimated with the kernel-linear (spline-flat) methodology. In contrast, VIX and central volatility are highly correlated 99.94% (99.98%). Moreover, RIX is only weakly correlated with the skewness and kurtosis showing that RIX captures left tail-risk differently than traditional risk-neutral moments.

2.6 Robustness tests

To address concerns that our results are driven by our particular choice of distribution, the analysis is repeated with two alternative distribution choices. First, a sinh-arcsinh transformed (Jones & Pewsey, 2009) normal distribution and second, a mixture of two normal distributions to capture multi-modal distributions. The sinh-arcsinh distribution transforms a random variable with a standard normal distribution to a new random variable with different skewness and kurtosis. The transformation is defined as follows:

$$Y_{\epsilon,\delta}(x) \equiv sinh[\delta sinh^{-1}(z) - \zeta] \quad z \sim N(0,1)$$

ζ controls the symmetry of the distribution of the new random variable $Y_{\epsilon,\delta}(z)$ and δ controls its tail weight. The hyperbolic sine function is denoted by $sinh$. A property of the transformation is that it defaults to the standard normal distribution if $\zeta = 0$ and $\delta = 1$. The advantage of the sinh-arcsinh transformation is that it has a parsimonious form that allows us to modify higher moments and offers easy implementation. Pricing European options based on the sinh-arcsinh distribution and the mixture of normals distribution is done by a Monte Carlo simulation with 100'000 repetitions and antithetic variables. The mean of the prices is forced to be equal to the forward price to make sure that our sample is arbitrage-free.

The results for the sinh-arcsinh distribution are presented in Tables 17 - 19 in the Appendix. The results mirror the previous results of the SVJ model. In Table 17 we can observe the same bias pattern as in the results of the SVJ generated option prices. Linear extrapolation leads to a faster convergence with lower initial errors than horizontal extrapolation. Moreover, model-free quantile moments have lower errors than their central counterparts, especially if the domain of available option prices is narrow. Table 18 shows that methods with flat extrapolation are less affected by micro-structural noise. Errors from different strike price spacings (Table 19) are small

compared to the bias induced by narrow domains. The observations are similar if the simulation study is based on a mixture of two normal distributions (Table 20 - 22 in the Appendix).

Table 16: Summary statistics of option-implied moments

This table displays descriptive statistics of quantile and central risk-neutral moments for the S&P 500. Quantile volatility is the inter-quartile range which is defined as the difference between the 75^{th} and 25^{th} percentile. Quantile skewness is defined by Hinkley's (1975) measure at the 90^{th} percentile. Quantile kurtosis is defined by Ruppert's (1987) ratio of quantile ranges with an outer range between the 95^{th} and 5^{th} percentile and an inner range between the 75^{th} and 25^{th} percentile. The central risk-neutral moments are based on Bakshi et al. (2003). Moments are estimated from a smoothed implied volatility surface based on a local-linear kernel regression with linear extrapolation or a cubic spline with flat extrapolation. Panel A shows descriptive statistics for the moments. Quantile and central kurtosis are reported as excess kurtosis over the normal distribution. Panel B shows the correlation among the time series of the moments estimated from the kernel-linear method. Panel C shows the correlation among the time series of the moments estimated from the spline-flat method. The sample spans the period between January 1996 and December 2017. All volatility measures, including VIX and SVIX, are displayed in percent. The values of RIX are multiplied by 1000.

Panel A: Descriptive statistics

	Kernel-linear					Spline-flat				
	Mean	*Std.*	*Med.*	*Min*	*Max*	*Mean*	*Std.*	*Med.*	*Min*	*Max*
Quantile Volatility	4.77	2.83	4.07	0.66	31.39	4.80	2.89	4.15	0.40	34.50
Quantile Skewness	-0.18	0.13	-0.19	-0.59	0.36	-0.21	0.12	-0.20	-0.62	0.18
Quantile ex. Kurtosis	0.52	0.64	0.35	-0.82	6.43	0.46	0.59	0.29	-0.91	7.13
Central Volatility	19.51	8.68	17.95	6.83	86.67	19.09	8.40	17.67	6.72	83.07
Central Skewness	-0.97	0.49	-0.93	-5.83	0.31	-0.79	0.37	-0.77	-2.29	0.32
Central ex. Kurtosis	1.93	2.72	1.38	-1.31	59.55	1.04	0.83	0.91	-1.08	6.85
VIX	19.34	8.47	17.87	6.83	82.56	18.98	8.27	17.56	6.71	80.08
RIX	1.50	5.28	0.48	0.01	165.90	1.13	3.15	0.42	0.01	81.16
SVIX	14.86	9.25	12.50	1.77	88.05	14.56	8.98	12.33	1.73	88.38

continued on the next page

Panel B: Correlation between moments - kernel-linear

		1	2	3	4	5	6	7	8	9
Quantile Volatility	1	100.00								
Quantile Skewness	2	-35.54	100.00							
Quantile Kurtosis	3	15.61	-42.30	100.00						
Central Volatility	4	75.05	-16.12	4.09	100.00					
Central Skewness	5	-24.62	60.10	-54.73	0.32	100.00				
Central Kurtosis	6	11.67	-28.14	38.37	0.60	-79.74	100.00			
VIX	7	74.21	-15.15	3.05	99.94	2.25	-1.02	100.00		
RIX	8	58.87	-19.03	12.53	59.80	-23.35	19.23	57.30	100.00	
SVIX	9	94.54	-45.31	43.81	69.74	-39.39	24.05	68.52	58.92	100.00

Panel C: Correlation between moments - spline-flat

		1	2	3	4	5	6	7	8	9
Quantile Volatility	1	100.00								
Quantile Skewness	2	-39.48	100.00							
Quantile Kurtosis	3	13.97	-60.26	100.00						
Central Volatility	4	71.40	-12.44	4.40	100.00					
Central Skewness	5	-7.19	73.84	-60.62	8.74	100.00				
Central Kurtosis	6	-13.08	-49.03	48.74	-22.22	-92.29	100.00			
VIX	7	70.86	-11.49	3.70	99.98	9.67	-22.94	100.00		
RIX	8	68.12	-25.10	12.66	67.14	-10.06	-1.74	66.11	100.00	
SVIX	9	94.24	-53.97	41.89	67.49	-22.74	-0.52	66.76	65.42	100.00

2.7 Conclusion

We contribute to the literature by performing an in-depth study on the robust estimation of risk-neutral moments. In theory, the estimation of risk-neutral moments is a straightforward task. Based on the theorem of Breeden & Litzenberger (1978) it is possible to obtain risk-neutral moments from the option-price surface. The caveat of this approach is that it assumes a continuum of option prices that spans a large domain. Real-world option data, however, is discrete and often spans only narrow domains with large gaps between strike prices. In addition, option-prices are not free from micro-structural noise introduced through bid-ask spreads and asynchronous trading of options. We test three different inter- and extrapolation techniques: cubic smoothing splines with horizontal extrapolation, cubic smoothing splines with linear extrapolation, and a local-linear kernel regression with linear extrapolation. These methods are benchmarked against a plain-vanilla implementation of Bakshi et al. (2003). Furthermore, we propose to use of quantile moments, which are more flexible than their central counterparts and allow for a more robust estimation.

Based on a SVJ model we simulate option data to test how strongly estimates of risk-neutral moments are affected by narrow domains, large strike price spacings, and micro-structural noise. Our findings suggest that estimates of risk-neutral moments are highly dependent on the estimation technique. Methods that rely on horizontal or no extrapolation tend to carry a large bias under narrow domains leading to a severe misestimation of risk-neutral moments. However, these methods are less affected by micro-structural noise. In contrast, methods that rely on linear extrapolation deliver more accurate estimates of risk-neutral moments under narrow domains, but are more affected by micro-structural noise. These results reveal a classic bias-variance trade-off between different estimation-methods. Two methods offer especially favorable bias-variance trade-offs. First, model-free quantile moments estimated from a local-linear kernel regression with linear extrapolation have only a small bias even under narrow domains and are only moderately affected by micro-structural noise.

Second, central risk-neutral moments based on a cubic smoothing spline with horizontal extrapolation are not strongly affected by micro-structural noise and their bias under narrow domains is acceptable. A viable alternative to cubic-smoothing splines are local-constant kernel regressions that carry a similar level of bias and variance. Moreover, if either a kernel regression or a smoothing spline is used, gaps between strike prices have a neglectable effect on the estimates of risk-neutral moments. We recommend to base the decision with respect to the estimator on the planned use of the risk-neutral moment estimates and the properties of the dataset. Model-free quantile moments obtained from a linearly extrapolated volatility surface should be used if a low bias in the estimates is required. In contrast, if researchers prefer a low variance of estimates, they should implement the central moments of Bakshi et al. (2003) and base them on a volatility surface interpolated with a cubic smoothing spline and extrapolated horizontally.

Moreover, we also analyze the sensitivity of the SVIX (Martin, 2017), the RIX (Gao, Gao, & Song, 2018; Gao, Lu, & Song, 2018), and the VIX. The results of risk neutral moments extend also to the SVIX, RIX, and VIX. Linear extrapolation leads to a higher sensitivity to micro-structural noise but reduces the bias compared to a flat extrapolation. However, it should be noted that SVIX and VIX are volatility indices and are thus much less affected by the extrapolation as prices of far-out-of-the-money options carry less weight compared to higher order moments. In contrast, RIX captures all higher order cumulants of the risk-neutral distribution and is thus a viable alternative for researcher that aim to obtain a left tail risk index.

The same results are also observable when comparing empirical risk-neutral moments of the S&P 500 index. Estimates based on a volatility surface obtained from cubic splines with horizontal extrapolation have lower standard deviations than estimates based on a volatility surface smoothed with a local-linear kernel regression and extrapolated linearly. Moreover, estimates of central skewness and kurtosis are higher

if they are based on the kernel-linear method, an effect likely due to a downward bias induced by horizontal extrapolation.

2.8 Appendix: Robustness tests

Table 17: Errors from truncated domain half-width - transformed normal distribution

This table shows the approximation errors of option-implied moments for different strike price domain half-widths in percent. The strike price spacing is held constant at 50 cents. The return distribution is a sinh-arcsinh transformed normal distribution. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.12 (0.20), skewness of -0.46 (-1.44), and kurtosis of 2.52 (8.34). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.18 (0.40), skewness of -0.63 (-2.4), and kurtosis of 3.17 (13.98). For each of the scenarios risk-neutral moments are estimated under different domain half-widths (10%, 50%, & 80%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Domain half-width	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	10.61	0.02	0.02	56.21	0.10	0.06	66.38	0.28	0.09
BKM-spline-flat	2.80	0.03	0.03	21.49	0.25	0.25	29.30	0.13	0.13
BKM-spline-linear	1.27	0.03	0.03	11.07	0.26	0.26	22.93	0.09	0.09
BKM-kernel-linear	1.08	0.00	0.00	9.53	0.07	0.07	19.97	0.33	0.33
BKM-lckernel-linear	2.63	0.00	0.00	20.46	0.00	0.00	27.19	0.02	0.02
Quantile Moments									
MFree-kernel-linear	0.25	0.24	0.24	1.88	0.58	0.58	0.81	0.18	0.18
MFree-spline-flat	0.22	0.20	0.20	15.74	2.33	2.33	3.25	0.60	0.60
MFree-lckernel-linear	2.14	1.47	1.47	17.07	3.93	3.93	0.72	1.47	1.47
BS-kernel-linear	7.29	7.29	7.29	24.72	25.49	25.49	1.40	3.30	3.30
BS-spline-flat	7.26	7.04	7.04	32.02	25.48	25.48	9.60	2.84	2.84
BS-lckernel-linear	7.87	7.88	7.88	30.88	24.32	24.32	9.93	3.41	3.41
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	32.83	2.46	0.02	84.88	14.84	0.42	91.34	33.03	1.35
BKM-spline-flat	14.91	1.22	0.02	55.07	8.17	0.18	67.92	20.15	0.57
BKM-spline-linear	4.06	0.17	0.01	11.43	1.37	0.05	18.22	3.88	0.07
BKM-Kernel-linear	3.25	0.11	0.01	9.29	1.01	0.13	14.59	2.90	0.30
BKM-lckernel-linear	15.13	1.35	0.07	55.89	9.63	1.02	69.79	24.58	3.86
Quantile Moments									
MFree-kernel-linear	0.34	0.07	0.07	3.27	0.01	0.01	0.89	0.02	0.02
MFree-spline-flat	14.32	0.23	0.45	3.46	0.54	2.36	8.10	0.14	0.23
MFree-lckernel-linear	16.26	1.08	0.79	2.75	0.64	0.61	10.57	0.61	0.33
BS-kernel-linear	34.30	33.97	33.97	21.67	24.14	24.14	8.73	13.59	13.59
BS-spline-flat	26.63	33.57	33.61	57.42	24.27	24.59	26.86	13.17	13.17
BS-lckernel-linear	25.68	33.12	33.12	57.25	23.24	23.24	26.33	13.27	13.27

Table 18: Errors from micro-structural noise - transformed normal distribution

This table shows the standard deviation of estimates of option-implied moments for different levels of micro-structural noise in percent. The strike price spacing set to 2.5 dollar and strikes cover a range from 90% to 110% of the strike price. The return distribution is a sinh-arcsinh transformed normal distribution. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.12 (0.20), skewness of -0.46 (-1.44), and kurtosis of 2.52 (8.34). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.18 (0.40), skewness of -0.63 (-2.4), and kurtosis of 3.17 (13.98). Micro-structure noise is simulated by perturbing option-prices by a percent of their value. Percentages are drawn randomly from a normal distribution with a standard deviation of either 1%, 5%, or 10%. For each level of micro-structural noise risk-neutral moments are estimated 1000 times. Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Noise	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.12	0.57	1.13	0.17	0.83	1.70	0.09	0.46	0.91
BKM-spline-flat	0.17	17.20	16.20	0.37	6.35	7.03	0.33	5.69	6.10
BKM-spline-linear	16.96	104.74	121.49	21.60	135.98	594.72	39.23	228.66	3.61e4
BKM-kernel-linear	2.60	7.28	14.95	9.17	19.99	27.10	16.38	35.05	45.82
BKM-lckernel-linear	0.15	0.98	1.60	0.47	2.64	4.74	0.60	3.42	5.75
Quantile Moments									
MFree-kernel-linear	7.36	9.30	11.65	4.31	8.64	9.95	7.71	11.21	15.71
MFree-spline-flat	5.64	21.72	26.05	2.10	31.71	29.73	5.72	247.89	91.03
MFree-lckernel-linear	5.09	13.94	16.13	7.78	22.85	28.51	5.91	15.83	20.29
BS-kernel-linear	0.46	1.49	2.95	2.05	5.28	8.54	3.96	10.79	20.28
BS-spline-flat	0.47	9.62	10.12	0.29	2.08	4.90	0.32	19.49	13.97
BS-lckernel-linear	0.32	1.74	3.26	0.49	1.97	3.66	0.34	1.75	3.35
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.15	0.79	1.55	0.41	2.06	4.00	0.39	1.99	3.85
BKM-spline-flat	0.16	12.70	29.49	0.48	17.85	31.54	0.58	6.17	12.22
BKM-spline-linear	0.19	70.71	118.76	1.07	155.66	267.23	2.44	370.04	575.00
BKM-Kernel-linear	0.21	0.96	1.96	1.35	5.57	14.90	3.00	12.66	45.85
BKM-lckernel-linear	0.16	0.99	1.71	0.37	1.92	3.76	0.31	1.70	3.25
Quantile Moments									
MFree-kernel-linear	3.58	8.19	10.49	4.83	12.06	16.81	3.74	9.57	12.77
MFree-spline-flat	1.05	16.08	23.94	1.41	26.73	41.71	1.21	1.55e3	1.14e3
MFree-lckernel-linear	7.59	13.99	18.27	5.49	10.86	14.75	8.79	16.45	22.73
BS-kernel-linear	0.30	1.30	2.44	0.32	1.17	2.10	0.30	1.43	2.76
BS-spline-flat	0.23	6.49	5.31	0.31	8.97	20.02	0.26	50.15	5.29
BS-lckernel-linear	0.44	2.38	3.74	0.32	1.43	2.55	0.31	1.71	2.66

Table 19: Errors from different strike price spacing - transformed normal distribution

This table shows the approximation error of option-implied moments for different strike price spacings in percent. The domain-half width is set to 99% of the current stock price. The return distribution is a sinh-arcsinh transformed normal distribution. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.12 (0.20), skewness of -0.46 (-1.44), and kurtosis of 2.52 (8.34). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.18 (0.40), skewness of -0.63 (-2.4), and kurtosis of 3.17 (13.98). For each of the scenarios risk-neutral moments are estimated under different strike price spacings (1%, 2%, & 5%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Spacing	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.09	0.18	0.27	0.21	0.43	8.46	0.36	0.71	2.88
BKM-spline-flat	0.03	0.04	0.06	0.21	0.27	0.46	0.13	0.09	0.18
BKM-spline-linear	0.03	0.04	0.05	0.22	0.28	0.49	0.09	0.06	0.09
BKM-kernel-linear	0.01	0.04	0.22	0.08	0.04	0.41	0.35	0.06	0.42
BKM-lckernel-linear	0.00	0.01	0.09	0.00	0.01	0.10	0.03	0.07	0.28
Quantile Moments									
MFree-kernel-linear	0.52	0.23	4.07	1.25	0.08	1.58	0.55	0.23	3.59
MFree-spline-flat	0.03	1.43	1.59	1.83	1.15	3.92	0.33	2.10	2.15
MFree-lckernel-linear	3.53	2.31	20.19	3.67	9.34	15.68	2.20	4.17	12.77
BS-kernel-linear	7.28	7.23	7.08	25.49	25.49	25.23	3.29	3.25	3.16
BS-spline-flat	7.05	7.18	7.09	25.49	25.77	25.45	2.84	2.93	2.93
BS-lckernel-linear	8.76	7.30	11.22	26.10	26.77	18.93	5.15	4.03	7.41
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.02	0.05	0.10	0.14	0.06	0.88	0.14	0.13	0.38
BKM-spline-flat	0.02	0.02	0.02	0.15	0.15	0.21	0.47	0.47	0.73
BKM-spline-linear	0.01	0.01	0.01	0.06	0.06	0.05	0.10	0.10	0.10
BKM-kernel-linear	0.01	0.01	0.04	0.11	0.09	0.03	0.21	0.20	0.02
BKM-lckernel-linear	0.03	0.02	0.01	0.44	0.41	0.29	1.63	1.55	1.29
Quantile Moments									
MFree-kernel-linear	0.08	0.22	0.14	0.18	1.17	0.20	0.20	0.06	0.73
MFree-spline-flat	0.32	0.27	1.02	2.24	2.38	0.90	0.16	0.14	0.81
MFree-lckernel-linear	2.1	6.70	6.02	4.52	4.01	19.92	1.11	3.32	15.52
BS-kernel-linear	33.96	33.95	33.81	24.15	24.18	24.40	13.58	13.56	13.42
BS-spline-flat	33.65	33.69	33.58	24.58	24.58	24.42	13.22	13.26	13.17
BS-lckernel-linear	32.3	33.93	30.77	23.69	24.60	24.13	12.83	12.82	13.33

Table 20: Errors from truncated domain half-width - mixture of 2 normal distributions

This table shows the approximation errors of option-implied moments for different strike price domain half-widths in percent. The strike price spacing is held constant at 50 cents. The return distribution is a multimodal mixture of 2 normal distributions. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.04 (0.20), skewness of -0.62 (-2.44), and kurtosis of 7.28 (13.25). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.09 (0.40), skewness of -0.15 (-3.57), and kurtosis of 6.16 (22.80). For each of the scenarios risk-neutral moments are estimated under different domain half-widths (10%, 50%, & 80%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Domain half-width	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	19.90	0.02	0.02	51.98	0.07	0.06	73.46	0.17	0.10
BKM-spline-flat	8.76	0.04	0.01	22.53	0.22	0.12	40.02	0.13	0.00
BKM-spline-linear	13.97	0.04	0.01	203.24	0.23	0.12	411.34	0.13	0.01
BKM-kernel-linear	12.82	0.00	0.00	48.70	0.03	0.03	182.52	0.03	0.04
BKM-lckernel-linear	8.33	0.00	0.00	21.47	0.01	0.01	38.69	0.00	0.00
Quantile Moments									
MFree-kernel-linear	0.26	0.25	0.25	4.69	2.06	2.06	10.52	0.07	0.07
MFree-spline-flat	8.29	10.27	6.59	10.52	5.50	5.40	15.92	7.76	4.31
MFree-lckernel-linear	0.60	2.68	4.45	6.71	5.23	8.08	7.02	2.16	3.99
BS-kernel-linear	177.38	177.38	177.38	15.51	19.36	19.36	51.12	57.44	57.44
BS-spline-flat	179.32	175.96	178.42	28.93	18.77	20.68	64.85	57.61	57.93
BS-lckernel-linear	175.71	175.65	175.90	29.41	19.49	19.77	64.52	57.56	57.65
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	45.94	5.76	0.04	82.65	17.03	0.08	92.10	35.10	0.87
BKM-spline-flat	32.73	2.32	0.01	59.96	7.42	0.04	75.76	17.37	0.33
BKM-spline-linear	50.00	1.09	0.00	50.67	4.21	0.02	80.52	10.72	0.20
BKM-kernel-linear	31.86	0.70	0.01	48.47	2.79	0.05	83.54	7.52	0.23
BKM-lckernel-linear	33.06	2.37	0.01	60.64	7.09	0.25	76.20	17.23	0.43
Quantile Moments									
MFree-kernel-linear	0.75	0.76	0.76	39.36	1.24	1.35	29.83	1.00	1.01
MFree-spline-flat	0.81	3.53	1.75	227.92	5.67	25.88	17.49	1.74	0.02
MFree-lckernel-linear	0.83	0.91	1.25	233.59	10.26	7.50	16.53	0.80	3.15
BS-kernel-linear	127.11	126.68	126.68	332.62	273.51	273.51	11.96	33.40	33.40
BS-spline-flat	114.13	125.71	124.79	113.57	273.72	272.33	61.78	33.34	32.88
BS-lckernel-linear	110.05	122.17	122.17	120.23	289.94	289.94	60.59	30.70	30.68

Table 21: Errors from micro-structural noise - mixture of 2 normal distributions

This table shows the standard deviation of estimates of option-implied moments for different levels of micro-structural noise in percent. The strike price spacing set to 2.5 dollar and strikes cover a range from 90% to 110% of the strike price. The return distribution is a multimodal mixture of 2 normal distributions. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.04 (0.20), skewness of -0.62 (-2.44), and kurtosis of 7.28 (13.25). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.09 (0.40), skewness of -0.15 (-3.57), and kurtosis of 6.16 (22.80). Micro-structure noise is simulated by perturbing option-prices by a percent of their value. Percentages are drawn randomly from a normal distribution with a standard deviation of either 1%, 5%, or 10%. For each level of micro-structural noise risk-neutral moments are estimated 1000 times. Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Noise	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.14	0.73	1.54	0.23	1.13	2.18	0.15	0.76	1.54
BKM-spline-flat	0.21	1.22	2.44	0.48	2.36	5.06	0.51	2.50	5.30
BKM-spline-linear	9.41	152.71	261.96	63.80	146.72	923.29	72.84	212.15	765.88
BKM-kernel-linear	4.51	26.14	58.22	16.80	61.81	79.48	49.04	163.85	188.92
BKM-lckernel-linear	0.15	0.76	1.55	0.37	1.91	3.70	0.54	2.84	5.36
Quantile Moments									
MFree-kernel-linear	5.98	23.46	32.28	3.50	16.71	28.33	6.45	26.25	29.32
MFree-spline-flat	3.79	23.73	47.62	1.01	9.39	30.02	3.20	21.44	46.30
MFree-lckernel-linear	18.22	26.90	31.42	11.08	21.30	28.39	21.87	26.69	27.41
BS-kernel-linear	1.22	5.95	9.86	1.98	9.66	16.59	2.54	14.99	32.06
BS-spline-flat	1.41	8.22	13.91	0.46	1.83	4.17	0.21	1.11	2.08
BS-lckernel-linear	0.96	7.25	11.38	0.24	1.24	2.73	0.41	1.25	1.94
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.09	0.48	0.96	0.12	0.62	1.22	0.04	0.22	0.43
BKM-spline-flat	0.20	1.01	2.17	0.35	1.75	3.65	0.19	0.99	2.01
BKM-spline-linear	38.90	139.04	143.77	18.99	93.98	2.84e3	34.59	1.07e4	3.43e3
BKM-Kernel-linear	12.50	45.59	71.27	3.76	16.09	26.23	3.63	25.59	38.26
BKM-lckernel-linear	0.17	0.88	1.77	0.32	1.56	3.10	0.26	1.28	2.51
Quantile Moments									
MFree-kernel-linear	8.39	13.99	19.05	52.30	104.75	140.93	10.90	18.74	21.51
MFree-spline-flat	2.73	24.14	32.20	5.10	53.05	107.38	2.24	69.86	238.43
MFree-lckernel-linear	9.49	18.08	21.96	89.22	136.16	186.31	9.30	18.88	30.65
BS-kernel-linear	3.20	28.90	179.64	23.04	64.50	89.00	7.81	20.67	24.29
BS-spline-flat	0.77	4.37	8.89	1.82	9.23	19.63	0.07	0.49	1.01
BS-lckernel-linear	5.62	7.05	9.68	3.50	8.27	13.53	1.29	1.76	2.52

Table 22: Errors from different strike price spacing - mixture of 2 normal distributions

This table shows the approximation error of option-implied moments for different strike price spacings in percent. The domain-half width is set to 99% of the current stock price. The return distribution is a multimodal mixture of 2 normal distributions. Option-implied moments are estimated under two different scenarios: Panel A shows the results of a standard scenario with quantile (central) volatility of 0.04 (0.20), skewness of -0.62 (-2.44), and kurtosis of 7.28 (13.25). Panel B displays the results of a crisis scenario with quantile (central) volatility of 0.09 (0.40), skewness of -0.15 (-3.57), and kurtosis of 6.16 (22.80). For each of the scenarios risk-neutral moments are estimated under different strike price spacings (1%, 2%, & 5%). Moments are estimated from raw option data (raw), or a smoothed volatility surface based on cubic splines with flat or linear extrapolation, a non-parametric local-linear kernel regression (kernel), or a local-constant kernel regression (lckernel). Quantile moments are estimated either model-free (MFree) or from a naïve Black-Scholes approximation (BS). The central risk-neutral moments are based on Bakshi et al. (2003).

	Volatility			**Skewness**			**Kurtosis**		
Spacing	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
Central Moments									
BKM-raw-none	0.09	0.18	0.79	0.24	0.46	1.52	0.37	0.72	2.34
BKM-spline-flat	0.01	0.02	0.02	0.13	0.16	0.17	0.02	0.07	0.10
BKM-spline-linear	0.01	0.02	0.02	0.13	0.16	0.17	0.03	0.07	0.10
BKM-kernel-linear	0.00	0.02	0.12	0.08	0.28	1.74	0.07	0.20	1.17
BKM-lckernel-linear	0.00	0.02	0.12	0.04	0.17	1.03	0.02	0.11	0.69
Quantile Moments									
MFree-kernel-linear	3.20	4.91	14.14	1.35	0.76	4.40	2.62	5.54	9.36
MFree-spline-flat	8.72	11.83	14.45	4.35	4.47	3.26	6.26	9.13	11.78
MFree-lckernel-linear	1.30	21.63	25.86	12.02	8.71	26.98	1.25	18.01	20.99
BS-kernel-linear	177.47	177.59	178.09	19.40	20.05	23.14	57.46	57.48	57.64
BS-spline-flat	177.98	178.59	179.58	20.34	19.89	19.58	57.85	58.06	58.35
BS-lckernel-linear	175.27	175.27	151.10	19.80	19.17	18.06	57.56	57.56	53.14
Panel B: Crisis Scenario									
Central Moments									
BKM-raw-none	0.01	0.05	0.12	0.02	0.19	0.21	0.36	0.09	0.79
BKM-spline-flat	0.02	0.02	0.03	0.06	0.06	0.08	0.33	0.32	0.40
BKM-spline-linear	0.02	0.02	0.04	0.05	0.04	0.04	0.28	0.27	0.26
BKM-kernel-linear	0.00	0.00	0.06	0.06	0.00	0.37	0.21	0.28	0.60
BKM-lckernel-linear	0.01	0.01	0.10	0.16	0.12	0.13	0.38	0.42	0.76
Quantile Moments									
MFree-kernel-linear	0.79	0.85	3.32	5.92	20.11	51.10	1.27	1.83	3.17
MFree-spline-flat	4.63	3.96	0.66	21.55	19.66	31.10	3.60	3.12	2.89
MFree-lckernel-linear	1.78	2.52	8.18	27.01	144.56	22.29	1.37	5.79	31.86
BS-kernel-linear	126.69	126.64	126.81	273.38	273.01	270.26	33.40	33.38	33.46
BS-spline-flat	125.09	125.19	125.66	272.98	272.59	274.63	33.15	33.16	33.21
BS-lckernel-linear	122.50	122.28	117.09	288.55	292.60	296.12	30.47	29.65	27.38

3 Volatility Control of Option Strategies

Alexander Feser

Status: Accept for publication in the Journal of Portfolio Management

ABSTRACT

Option trading strategies can be managed by taking less risk if volatility is high and more risk if volatility is low. These volatility managed option strategies generate economically and statistically significant alphas over their unmanaged counterparts, have reduced maximum drawdowns, lower downside risk, and more normal return distributions. The findings hold for 9 out of 10 of the S&P500 Cboe option strategy benchmarks and are especially strong for put- and buy-write strategies. Volatility controlling the put-write benchmark generates a significant annualized alpha of 4.88% per year over the unmanaged strategy. To understand the driver behind the success of volatility control, option strategy returns are decomposed into a return driven by changes in the S&P500 and a return which captures the exposure to the variance risk premium (VRP). The results show that the success of volatility control is driven by both timing the exposure to the VRP and the exposure to the S&P500.

JEL Classifications: G10, G11, G12

Keywords: volatility control, covered call, put-write, options

I thank Manuel Ammann, Stephan Kessler, Jan-Philip Schade, Nic Schaub, Paul Söderlind, Michael Verhofen, Felix von Meyerinck, and an anonymous referee for helpful discussions and comments. All errors are my own.

3.1 Introduction

Option trading strategies, such as put- and buy-write strategies, are popular investments. They offer returns comparable to an investment into the underlying at a lower volatility thus generating high Sharpe ratios (Israelov & Nielsen, 2015a). However, put- and buy-write strategies also carry high levels of downside risk which materializes in occasional crashes. This study shows that option strategies can be managed by reducing exposure when volatility is high and increasing exposure if volatility is low. These volatility controlled option trading strategies generate significant alphas over their unmanaged counterparts, have increased Sharpe ratios, lower absolute skewness, lower kurtosis, and less downside risk.

I argue that option strategies, especially put- and buy-write strategies, are natural candidates for volatility control which seems counterintuitive at first. Put-write strategies generate profits by collecting option premia and thus it is commonly expected that they are more profitable during times of high (implied) volatility as option premia are also higher. However, during these times they are also exposed to large losses as the underlying market is more volatile. The return distributions of put- and buy-write strategies consist of frequent small profits which are offset by occasional large losses. Volatility control helps to reduce these large losses and leads to significant increases in Sharpe ratios.

The results show that volatility controlled option trading strategies generate significant and positive alphas for 9 of 10 of the analyzed strategies. The results are particularly strong for put- and buy-write strategies, such as covered calls, where annualized alphas range from 3.45% to 4.88% and Sharpe ratios increase between 0.22 and 0.37. Moreover, volatility control reduces absolute skewness, kurtosis, and the drawdown of these strategies. This risk reduction makes the risk-return profile of these strategies more appealing for downside risk averse investors and challenges the findings of studies which explain the return of put-write strategies with their down-

side risk (e.g. Lettau et al., 2014). The results also extend to other option strategies such as collar, protective put, and condor strategies.

Furthermore, I also analyze the drivers of volatility control in option strategies. Option trading strategies expose investors to the equity risk premium and the variance risk premium (Coval & Shumway, 2001). Moreira & Muir (2017) show that volatility control works well for equity market returns and therefore a natural hypothesis is to assume that the success volatility control in option strategies is driven by controlling the exposure to the S&P500. Using the decomposition of Israelov & Nielsen (2015a) I investigate if the success of volatility control in option trading strategies is driven by controlling the strategy's exposure to the equity market or by controlling the exposure to the variance risk premium. Israelov & Nielsen (2015a) show that option strategy returns can be decomposed into (1) a return driven by changes in the underlying, (2) a return from the exposure to the volatility risk premium, and (3) a dynamic equity return originating from changes in the delta of options over time. The results show that the gains from volatility control are the most statistically significant for the variance risk premium return. However, alphas from controlling the equity risk premium are larger in magnitude, i.e. economically more significant, but do not have the same level of statistical significance. Therefore, the results are not driven by simply controlling the S&P500 but instead by jointly controlling the exposure to the variance and equity risk premium.

To the best of my knowledge, this study is the first to show that volatility control and the low volatility anomaly extend to option strategies. In particular, the results indicate that the increased risk-return trade-off is mainly driven by managing the exposure to the volatility risk premium and thus the results are distinct from previous studies which find that positions in the equity market can be managed with volatility control. The findings extend the low volatility anomaly to option and volatility investments and are particularly interesting because they contradict the common in-

tuition of practitioners (see the examples in Moreira & Muir, 2017) and academics (e.g. Bakshi & Kapadia, 2003, find that the variance risk premium is higher in times of high volatility) that investors should buy in times of high volatility.

This study is related to the literature on volatility control and option strategies. The economic value of volatility control is first discussed in depth by Fleming, Kirby, & Ostdiek (2001, 2003) who link it to a mean-variance framework. If investors expect returns to be constant, a mean-variance approach will lead to volatility control in the portfolio. Moreira & Muir (2017) follow this argument and show that volatility control generates large and significant alphas for risk-factors which lead to sizable utility gains for investors. The mechanism behind volatility control is described by Perchet, de Carvalho, Heckel, & Moulin (2016) who show that volatility control works because of three phenomena frequently observed in financial markets. The first phenomenon is volatility clustering. If the volatility in a period is high, it will likely remain high in the next period. Therefore, if exposure is reduced during times of high volatility, the overall volatility of the strategy is reduced leading to a higher Sharpe ratio. Volatility clustering can be measured as the autocorrelation of volatility. Second, periods of high volatility also often coincide with low returns. This negative correlation between returns and volatility is often referred to as leverage effect. If volatility correlates negatively with returns, reducing exposure in times of high volatility also helps to avoid downside returns in the strategy. Third, fat tails in the return distribution increase the efficacy of volatility control, typically by increasing the leverage effect. Economically, these phenomena can be tied to the financial leverage in companies balance sheets, business cycle risk, and self-exciting behavior of investors (Carr & Wu, 2017). Perchet et al. (2016) point out that even if returns follow a white noise process and volatility does not cluster, volatility control will not hurt the risk-return trade-off net of transaction costs.

This study proceeds as follows. Section 3.2 reviews the option trading strategies

in the sample and presents summary statistics. Section 3.3 presents the volatility control methodology and its effect on the risk-return trade-off. Section 3.4 discusses the effect of volatility control on risk. Section 3.5 investigates the drivers of the success of volatility control in option strategies. Section 3.6 shows that volatility control in option strategies is robust to transaction costs. Section 3.7 concludes.

3.2 Data

I obtain daily and monthly data of option strategy benchmarks from the Chicago Board Options Exchange (Cboe). I use all 10 published option strategy benchmarks covering the period from January 1990 to August 2017. Cboe also provides data for the S&P500 and the S&P500 Total Return Index. In addition, S&P500 option price data, are obtained from the OptionMetrics IvyDB through WRDS. S&P500 cash dividends are provided by Compustat. The yield on 3 month T-bills is sourced from the Federal Reserve of St Louis. The remainder of this section briefly describes the Cboe option strategy benchmarks and their summary statistics (Table 23). All results in this study are computed from excess returns over the 3 month T-bill yield.

The data set includes 10 option strategies, all of which are based on the S&P 500. Cboe publishes these strategies as benchmarks for investors and their data is widely used by other studies (e.g. Whaley, 2002; Ungar & Moran, 2009; Israelov & Nielsen, 2015a). Option positions in all strategies are rolled at expiration, typically on the third Friday of each month. The benchmark strategies typically choose the options with a targeted strike, e.g. the current level of the S&P 500 for an at-the-money (ATM) option. If no such option is available, the Cboe strategy benchmarks typically choose the next closest option which is further out-of-the-money (OTM) than the targeted strike. The strategies generally buy (sell) options at the last ask (bid) quote before 11 am. Therefore, the benchmark strategies include transaction costs arising from the bid-ask spread but ignore any further transaction costs such as brokering

fees. Furthermore, the strategy benchmarks are fully collaterized and hold cash to cover losses from selling options. For example, the Cboe put-write benchmark holds cash equal to the strike price of the options sold and thus has a relatively low leverage. This is important because collaterized option strategy returns have relatively normal return distributions and can be evaluated with traditional statistical inference techniques. In contrast, uncollaterized option strategies have excessive fat tails which potentially imply that the moments of the return distribution do not exist anymore (Eraker, 2013) and consequently traditional inferential statistics cannot be applied anymore. The strategies can be roughly grouped into 3 categories: (1) put- and buy-write strategies such as covered calls (2) defensive strategies such as protective puts, and (3) volatility carry strategies that aim to pick up the variance risk premium.

The put- and buy-write category contains 3 covered call strategies, 1 put-write strategy and 1 combo strategy. All of these strategies have a concave payoff profile which is similar to a short put. Covered call strategies achieve this profile by combining a long position in the underlying with a short call. Thus, an investor collects the premium of the call option which provides a steady flow of returns but also forfeits all upside potential above the strike of the call while retaining the downside risk (see Whaley, 2002; Israelov & Nielsen, 2014, 2015a, for an in-depth discussion). The three covered call strategies in the sample sell at-the-money calls (BXM index), 2% out-of-the-money calls (BXY index), and OTM calls with a delta of 0.3 (BXMD index). The main difference between them is thus the strike price of the call option that is sold, e.g. the BXMD index typically sells call options with the highest strike prices. Thus, the option premia collected by the BXMD index are the lowest of the three covered call strategies, but it also retains the most upside potential. The put-write strategy (PUT index) sells ATM puts and fully collaterizes them. For example, if a put option with a strike of 100 is sold, the strategy will invest 100 in a money market account to cover all potential losses. Therefore, the PUT index generates the same payoff profile as the BXM covered call strategy. Furthermore, the sample contains a combo

strategy (CMBO index) which combines writing ATM puts and a 2% OTM covered call strategy. In comparison to the BXM index it adds extra downside exposure if the price of the S&P500 declines.

The concave payoff profile of the put- and buy-write strategies is also reflected in their summary statistics. Covered call, put-write and combo strategies have similar return distributions. All five strategies have higher returns and lower volatilities and therefore higher Sharpe ratios (Sharpe, 1994) than the S&P500 index. This outperformance over the S&P500 can also be seen in the significantly positive CAPM alphas of the five strategies. The most common explanations for this outperformance are either that selling options collects the variance risk premium (Coval & Shumway, 2001; Israelov & Nielsen, 2015b) or a downside risk premium (Lettau et al., 2014). Indeed, all five strategies have left skewed and fat tailed return distributions in comparison with the S&P 500 indicating that the strategies have a high downside risk. This observation is also confirmed by comparing the CAPM beta (β) with the downside beta (β^-) of Lettau et al. (2014). The downside beta is a conditional beta calculated from returns when the stock market is in distress.[8] Therefore, the downside beta provides a measure of systematic downside risk of an investment. The summary statistics show that all five put- and buy-write strategies have a higher downside beta than their regular CAPM beta. However, the downside beta is not higher than 1 as the strategies just forfeit the upside during bull markets but retain the downside during bear markets. Therefore, they have betas close to 1 if the market declines. Moreover, put- and buy-write strategies have lower drawdowns than the S&P500 indicating that these strategies have actually a lower downside risk than the S&P500. All put- and buy-write strategies exhibit a strong negative correlation between returns and volatility as well as a high autocorrelation in volatility which makes them ideal candidates for volatility control. Interestingly, the PUT index has a higher Sharpe ratio than the BXM index. Both strategies use at-the-money options and put-call parity de-

[8] Lettau et al. (2014) define a downside event as a return that is smaller than the average return minus one standard deviation.

mands that they should have an identical payoff profile. Israelov (2017) analyzes this discrepancy and attributes it to different exposures to the S&P500 during expiration days.

The second category of defensive strategies includes three option strategies: a collar, a zero-cost collar, and a protective put strategy. The protective put strategy (PPUT index) limits the downside of an S&P 500 investment by buying 5% OTM puts. This protective put strategy is a classic hedge against large losses. However, the premia of the puts lower the return if no downside event takes place. The collar strategy (CLL index) combines a protective put (5% OTM) with a short 10% OTM call. Thus, downside and upside are both capped by the options. The zero-cost put-spread collar (CLLZ index) buys a 2.5% OTM put, sells a 5% OTM put (this combination of put options is referred to as a put-spread), and sells an OTM call to cover the cost of the put-spread. In addition, CLLZ buys one unit of the S&P500 to collateralize the short call. Therefore, it has a similar payoff profile as a collar, but provides a limited amount of downside-risk protection by limiting losses only between 2.5% and 5%. All three strategies have Sharpe ratios which are close to the Sharpe ratio of the S&P500 and thus do not enhance the risk-return trade-off. Moreover, they have on average lower returns and do not generate any alpha over the S&P500 index. However, the collar and protective put strategy have lower downside risk than the S&P500, indicated by slightly skewed returns, low excess kurtosis, and low downside betas. Protective put and collar strategies exhibit a high autocorrelation in realized volatility which points to volatility clustering. However, the leverage effect of the PPUT and CLL strategies is relatively low with correlations of -15.41% and -7.27% respectively.

The third category of strategies contains two volatility strategies: a butterfly (BFLY index) and a condor strategy (CNDR index). A butterfly strategy sells an ATM call and an ATM put (referred to as a short straddle) and limits the downside

risk by buying 5% OTM call and put options. Consequently, the profit of the butterfly is highest if the price of the underlying remains unchanged. However, a butterfly strategy is not a pure volatility investment. The profits of the strategy depend on the combined premium that is generated from selling the ATM options and buying the OTM options. If the volatility surface is exhibits a steep smile, i.e. OTM puts are relatively expensive, the profits of this strategy are lower. A variation of the butterfly is the condor strategy which sells an OTM call and an OTM put with deltas of +/- 0.2 (also referred to as a short strangle) and limits the risk by buying deep OTM calls and puts with deltas of +/- 0.05. Both strategies have low betas of 0.10 and 0.15, low downside betas of 0.38 and 0.46, and similar alphas over the S&P500 of 2.99% and 3.08% per year. However, the two strategies have different distributional characteristics. The BFLY strategy has a relatively normal distribution and a similar Sharpe ratio to the S&P500, while the condor strategy has a left skewed leptokurtic payoff profile. Both strategies exhibit a relatively low autocorrelation in their respective volatilities. In addition, the BFLY strategy has a low correlation of -10.16% between returns and volatility. It is worthwhile to note that both strategies have stopped performing after the financial crisis of 2008 and that most of the observed high Sharpe ratios are driven by a strong performance in the 1990's.

3.3 Volatility control

I construct volatility controlled versions of the option trading strategies similar to Moreira & Muir (2017) by scaling the option strategy returns by the inverse of their realized volatility. At the end of each month, the realized volatility determines the exposure to the option trading strategy over the following month. Thus, the return of the volatility controlled option strategy is:

$$\frac{c}{\sigma_t} r_{t,t+1} \tag{28}$$

where $r_{t,t+1}$ is the excess return of a option strategy in the subsequent month. The position in a option trading strategy is thus inversely determined by volatility (σ_t). If

Table 23: Summary statistics of option trading strategies

This table shows the summary statistics of the 10 option trading strategies in the sample and the S&P500 index. The sample period begins January 1990 and ends in August 2017 covering a total of 332 months. Column 1 shows the average annualized excess return over the 3 month T-bill yield, column 2 displays the annualized volatility, column 3 shows the skewness, column 4 shows the excess kurtosis, column 5 shows the annualized Sharpe ratio (SR), columns 6 shows the maximum drawdown (DD) during the sample period, column 7 shows the CAPM beta, column 8 shows the downside beta according to Lettau et al. (2014), column 9 shows the CAPM alpha, column 10 shows the correlation between the volatility and returns of the option strategy, and column 11 shows the autocorrelation of the volatility in the option strategy. The data covers 332 months per strategy during the period from January 1990 to August 2017. Standard errors of alphas are Newey & West (1987) adjusted with 5 lags and are unreported to keep the table concise.

	$\overline{ret}$	σ	$Skew$	$Kurt$	SR	DD	β	β^-	α	$\rho_{ret,\sigma}$	$\rho_{\sigma_t,\sigma_{t-1}}$
Cov. Call ATM (BXM)	6.30	10.09	-1.27	4.96	0.60	36.94	0.62	0.96	2.80 ***	-40.74	79.58
Cov. Call 2% OTM (BXY)	7.50	11.86	-0.88	2.81	0.61	41.49	0.78	0.94	3.14***	-38.88	80.74
Cov. Call 30Δ (BXMD)	8.01	12.34	-0.82	2.55	0.63	43.78	0.82	0.99	3.37***	-34.54	83.66
Put-Write (PUT)	7.45	9.65	-1.89	9.24	0.75	33.07	0.56	1.01	4.29***	-42.56	78.63
Covered Combo (CMBO)	6.79	10.49	-1.25	4.58	0.63	39.42	0.66	0.97	3.07***	-40.65	80.06
Collar (CLL)	3.71	10.25	-0.18	-0.22	0.36	37.93	0.66	0.18	0.11	-7.27	81.03
Zero-Cost Collar (CLLZ)	4.24	11.09	-0.92	2.78	0.37	51.01	0.74	0.91	0.16	-35.09	85.07
Protective Put (PPUT)	4.07	11.68	-0.34	0.30	0.34	49.30	0.75	0.19	0.00	-15.41	87.49
Butterfly (BFLY)	3.54	10.71	0.06	-0.37	0.32	34.00	0.10	0.38	2.99	-10.16	46.79
Condor (CNDR)	3.90	6.81	-2.11	5.90	0.56	16.60	0.15	0.46	3.08***	-42.71	48.47
S&P500 Index	5.50	14.26	-0.59	1.31	0.38	61.80	1.00	1.00	0.00	-29.76	87.21

Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$

volatility is low, the volatility controlled portfolio increases the position in a strategy and vice versa. The constant c allows to control the leverage used in the strategy and can be interpreted as the desired ex-ante volatility. However, as c is constant over time, it does not change the Sharpe ratio of the strategy. To make results between the volatility managed strategy and the unmanaged strategy comparable, I follow Moreira & Muir (2017) and set c ex-post so that the unconditional volatility of the managed strategy is the same as the unconditional volatility of the unmanaged strategy. This an easy way to make statistics between the managed and unmanaged option trading strategies comparable as they have the same amount of unconditional risk.

Realized volatility is computed from squared daily returns as an exponentially weighted moving average with the center of mass at 20 days, similarly to Moskowitz, Ooi, & Pedersen (2012). I follow Moskowitz et al. (2012) and Moreira & Muir (2017) and evaluate the success of volatility control by calculating the alpha of the volatility controlled strategy over the unmanaged strategy. I run the following regression to determine if volatility controlled strategies outperform an unmanaged strategy:

$$r_t^{managed} = \alpha + \beta r_t^{unmanaged} + \epsilon_t \tag{29}$$

Tables 24 and 25 present summary statistics of the volatility controlled strategies and the residual returns from the regression in Equation 29. The main results are presented in Table 26. In addition to the alpha, Table 26 also reports the regression beta, the Sharpe ratio of the managed strategies, the change in Sharpe ratio, and the information ratio. The information ratio, given by $\frac{\alpha}{\sigma(\epsilon)}$, measures by how much volatility control expands the risk-return trade-off. Therefore, it is often also referred to as the excess Sharpe ratio (Moreira & Muir, 2017). In addition, to create a better overview for the reader, Table 24 and 25 display summary statistics for the volatility controlled strategies and the regression residuals.

Table 24: Summary statistics of volatility controlled option trading strategies

This table shows the summary statistics of the 10 volatility controlled option trading strategies in the sample and the S&P500 index. The sample period begins January 1990 and ends in August 2017 covering a total of 332 months. Column 1 shows the average annualized excess return over the 3 month T-bill yield, column 2 displays the annualized volatility, column 3 shows the skewness, column 4 shows the excess kurtosis, column 5 shows the annualized Sharpe ratio (SR), columns 6 shows the maximum drawdown (DD) during the sample period, column 7 shows the CAPM beta, column 8 shows the downside beta according to Lettau et al. (2014), column 9 shows the CAPM alpha. Standard errors of alphas are Newey & West (1987) adjusted with 5 lags and are unreported to keep the table concise.

	$\overline{ret}$	σ	$Skew$	$Kurt$	SR	DD	β	β^-	α_{CAPM}
Cov. Call ATM (BXM)	9.35	10.09	-0.97	1.87	0.89	-29.18	0.56	0.33	5.97^{***}
Cov. Call 2% OTM (BXY)	10.59	11.86	-0.58	0.62	0.85	-36.07	0.71	0.31	6.28^{***}
Cov. Call 30Δ (BXMD)	11.34	12.34	-0.53	0.79	0.87	-35.34	0.74	0.32	6.80^{***}
Put-Write (PUT)	11.47	9.65	-1.50	4.78	1.13	-29.24	0.50	0.41	8.38^{***}
Covered Combo (CMBO)	10.02	10.49	-0.89	1.56	0.91	-31.80	0.61	0.36	6.34^{***}
Collar (CLL)	4.94	10.25	-0.02	-0.06	0.47	-31.29	0.62	0.04	1.33
Zero-Cost Collar (CLLZ)	6.63	11.09	-0.44	0.67	0.58	-38.43	0.67	0.31	2.65^{**}
Protective Put (PPUT)	5.69	11.68	-0.15	0.06	0.48	-36.28	0.69	-0.01	1.63
Butterfly (BFLY)	3.33	10.71	-0.12	-0.16	0.31	-35.03	0.07	0.31	2.91
Condor (CNDR)	4.93	6.81	-1.84	5.95	0.71	-15.53	0.11	0.23	4.30^{***}
S&P500 Index	7.92	14.26	-0.35	0.32	0.54	-44.92	0.91	0.30	2.55^{*}

*Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$*

Table 25: Summary statistics of residuals

This table shows the summary statistics of the residuals from the regression in Equation 29 for the 10 option trading strategies in the sample and the S&P500 index. The sample period begins January 1990 and ends in August 2017 covering a total of 332 months. Column 1 shows the average annualized excess return over the 3 month T-bill yield, column 2 displays the annualized volatility, column 3 shows the skewness, column 4 shows the excess kurtosis, column 5 shows the annualized Sharpe ratio (SR), columns 6 shows the maximum drawdown (DD) during the sample period, column 7 shows the CAPM beta, column 8 shows the downside beta according to Lettau et al. (2014), column 9 shows the CAPM alpha. Standard errors of alphas are Newey & West (1987) adjusted with 5 lags and are unreported to keep the table concise.

	$\overline{ret}$	σ	$Skew$	$Kurt$	SR	DD	β	β^-	α_{CAPM}
Cov. Call ATM (BXM)	0.00	5.21	0.79	5.54	0.00	-38.15	0.03	-0.49	-0.14
Cov. Call 2% OTM (BXY)	0.00	5.42	0.98	5.42	0.00	-38.59	0.01	-0.53	-0.08
Cov. Call 30Δ (BXMD)	0.00	6.05	0.81	3.93	0.00	-44.56	0.02	-0.55	-0.13
Put-Write (PUT)	0.00	5.34	0.82	6.36	0.00	-44.64	0.03	-0.44	-0.19
Covered Combo (CMBO)	0.00	5.20	0.92	5.32	0.00	-41.76	0.03	-0.48	-0.15
Collar (CLL)	0.00	2.99	0.60	2.17	0.00	-22.37	-0.01	-0.14	0.06
Zero-Cost Collar (CLLZ)	0.00	5.31	1.02	5.19	0.00	-34.37	0.02	-0.49	-0.11
Protective Put (PPUT)	0.00	3.90	0.91	4.24	0.00	-24.78	-0.01	-0.19	0.06
Butterfly (BFLY)	0.00	2.70	-0.47	5.19	0.00	-10.43	-0.02	-0.06	0.11
Condor (CNDR)	0.00	2.91	0.39	5.72	0.00	-19.02	-0.03	-0.19	0.16
S&P500 Index	0.00	6.12	0.76	4.07	0.00	-34.60	0.00	-0.61	0.00

Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$

Table 26: Performance of volatility managed option strategies

This table shows the performance of volatility managed trading strategies. The table reports the the alphas over the unmanaged strategy (α), the beta from the regression (β), the information ratio (IR), the Sharpe ratio of the managed strategies (SR), and the change in Sharpe ratio (ΔSR). The volatility controlled strategies are scaled to have the same unconditional volatility as the unmanaged strategies. The data covers 332 months per strategy during the period from January 1990 to August 2017. Standard errors are reported in parenthesis. Standard errors of alphas are Newey & West (1987) adjusted with 5 lags. Standard errors of ΔSR are obtained with a circular block bootstrap with 10'000 repetitions. All statistics are annualized. Alphas are reported in percent.

	α		β		IR	SR	ΔSR	
Cov. Call ATM (BXM)	3.64^{**}	(1.50)	0.86^{***}	(0.10)	0.69	0.89	0.26^{*}	(0.15)
Cov. Call 2% OTM (BXY)	3.45^{**}	(1.50)	0.89^{***}	(0.08)	0.63	0.85	0.22^{*}	(0.12)
Cov. Call 30Δ (BXMD)	3.88^{**}	(1.65)	0.88^{***}	(0.09)	0.63	0.87	0.23^{*}	(0.13)
Put-Write (PUT)	4.88^{***}	(1.78)	0.83^{***}	(0.11)	0.90	1.13	0.37^{**}	(0.19)
Covered Combo (CMBO)	3.74^{**}	(1.53)	0.87^{***}	(0.09)	0.71	0.91	0.27^{*}	(0.15)
Collar (CLL)	1.13^{*}	(0.58)	0.96^{***}	(0.04)	0.38	0.47	0.09	(0.06)
Zero-Cost Collar (CLLZ)	2.59^{*}	(1.35)	0.88^{***}	(0.10)	0.48	0.58	0.18	(0.13)
Protective Put (PPUT)	1.51^{*}	(0.78)	0.95^{***}	(0.05)	0.38	0.48	0.11	(0.07)
Butterfly (BFLY)	-0.19	(-0.48)	0.97^{***}	(0.03)	-0.07	0.31	-0.03	(-0.04)
Condor (CNDR)	1.37^{**}	(0.69)	0.90^{***}	(0.06)	0.47	0.71	0.15	(0.10)
S&P500 Index	2.55^{*}	(1.43)	0.91^{***}	(0.07)	0.41	0.54	0.14	(0.11)

Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$

The results in Table 26 show that volatility control delivers significantly positive alphas for 9 out of the 10 option trading strategies. The covered call and put-write strategies benefit especially from volatility control. A volatility managed BXM strategy generates an annualized alpha of 3.64% over its unmanaged counterpart which is statistically significant at the 5% level. The results are similar for the other put- and buy-write strategies which have statistically significant alphas of 3.45%, 3.88%, 4.88%, and 3.74% for the BXY, BXMD, PUT, and CMBO index respectively. These high alphas also translate into high information ratios and significant increases in Sharpe ratios. This finding is not unexpected due to the high autocorrelation of volatility and the strong leverage effect of put- and buy-write strategies. All these strategies collect the variance risk premium during calm times which leads to a steady stream of small returns with occasional large losses when the S&P500 decreases. The volatility control strategy reduces exposure in times of high volatility which coincide with down movements of the market.

The results also extend to option strategies which offer downside protection, such as collars and protective puts. However, the benefit of volatility control is smaller for these strategies as they have an inherently lower exposure to negative returns of the underlying and thus benefit to a lesser extent from the volatility reduction offered by volatility control during volatile market regimes. For example, the CLL collar strategy generates an annualized alpha of 1.13% which is statistically significant at the 10% level. However, the value of volatility control is small in comparison to volatility control the put-write strategy (alpha of 4.88%) or even the S&P500 (alpha of 2.55%). Moreover, volatility control delivers a 0.09 increase in Sharpe ratio. The effect of volatility control is of similar magnitude for the PPUT strategy generating an alpha of 1.51% and increasing the Sharpe ratio by 0.11. The findings are in line with the conclusions of Perchet et al. (2016). PPUT and CLL strategies have relatively low leverage effects and thus do not benefit from volatility control as much as put- and buy-write indices. Notably, the effect is stronger for the zero-cost collar strategy. The zero-cost collar strategy only protects against downside returns up to 5% and is fully exposed to downside events exceeding 5%. This is reflected in the significant alpha of 2.59% of the volatility controlled strategy over its unmanaged counterpart.

In contrast, volatility control is less effective for the butterfly strategy. Volatility control has no effect on the butterfly strategy in the sample and leads to an insignificant alpha of -0.19%. The relative ineffectiveness of volatility control for the butterfly strategy could be driven by the relatively small autocorrelation of the butterfly volatility combined with a small leverage effect and a relatively normal return distribution. Nevertheless, while volatility control does not increase the risk-return trade-off of the butterfly strategy, it also does no harm it. A volatility timed condor strategy on the other hand generates a significant alpha of 1.37% over its unmanaged counterpart and its Sharpe ratio increases by 0.15.

Furthermore, the findings from Table 26 can also be observed in Table 24. Note,

that volatility controlled and uncontrolled strategies have exactly the same volatility by construction. However, except for the BFLY strategy, all volatility controlled option strategies have higher returns than their unmanaged counterparts, resulting in increased Sharpe ratios. Volatility control also leads to a reduction of risk which is further discussed in Section 3.4. Summary statistics of the residuals in Table 25 show that the residuals have fat tails indicated by high values of excess kurtosis for all 10 strategies.

3.4 The impact of volatility control on risk

The impact of volatility control on the systematic and idiosyncratic risk of option trading strategies is shown in Table 27 and 28 respectively. To begin with, the results show that volatility control lowers the (absolute) skewness, kurtosis, and the maximum drawdown of the strategies. The risk-reduction is greatest among the put- and buy-write strategies. For example, the skewness of the BXM increases by 0.31 which reduces the negative skewness of the BXM returns to -0.97. Moreover, the kurtosis of the strategy decreases by -3.21 and the maximum drawdown is reduced by 7.76 percentage points. In addition, volatility control also reduces the CAPM beta of the BXM strategy by 0.06 and its downside beta by 0.62. Therefore, volatility control especially helps to reduce downside risk. An interesting observation is that the downside beta of the volatility controlled strategy is lower than the CAPM beta. Lettau et al. (2014) show that the returns of put-write strategies are explained by the downside risk that these strategies carry. However, the results of this study indicate that the volatility managed strategies defy this explanation. The results for the other four put- and buy-write strategies are similar. The absolute skewness, kurtosis, maximum drawdown, beta, and downside beta of put- and buy-write strategies is reduced. Volatility control thus does not only generate alpha, but also reduces the downside risk of the strategy.

Furthermore, the results extend to collar and protective put strategies although

Table 27: Systematic risk of volatility managed option strategies

This table shows the impact of volatility control on the systematic risk of option trading strategies. The table displays the CAPM beta and the downside beta of volatility managed option trading strategies, as well as their differences to the unmanaged strategies ($\Delta\beta$ and $\Delta\beta^-$). The data covers 332 months per strategy during the period from January 1990 to August 2017. Standard errors of $\Delta\beta$ and $\Delta\beta^-$ are obtained using a circular block bootstrap with 10'000 repetitions.

	β	β^-	$\Delta\beta$		$\Delta\beta^-$	
Covered Call ATM (BXM)	0.56	0.33	-0.06	(-0.07)	-0.62***	(-0.22)
Covered Call 2% OTM (BXY)	0.71	0.31	-0.07	(-0.07)	-0.63***	(-0.21)
Covered Call 30Δ (BXMD)	0.74	0.32	-0.08	(-0.08)	-0.67***	(-0.22)
Put-Write (PUT)	0.50	0.41	-0.06	(-0.06)	-0.60**	(-0.29)
Covered Combo (CMBO)	0.61	0.36	-0.06	(-0.07)	-0.61***	(-0.21)
Collar (CLL)	0.62	0.04	-0.04	(-0.03)	-0.14**	(-0.07)
Zero-Cost Collar (CLLZ)	0.67	0.31	-0.07	(-0.08)	-0.60***	(-0.22)
Protective Put (PPUT)	0.69	-0.01	-0.05	(-0.04)	-0.20	(-0.14)
Butterfly (BFLY)	0.07	0.31	-0.03**	(-0.01)	-0.08	(-0.05)
Condor (CNDR)	0.11	0.23	-0.04**	(-0.02)	-0.23**	(-0.11)
S&P500 Index	0.91	0.30	-0.09	(-0.08)	-0.70***	(-0.24)

*Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$*

Table 28: Idiosyncratic risk of realized volatility managed portfolios

This table shows the impact of volatility control on the risk of option trading strategies. The table displays the skewness ($skew$), the excess kurtosis ($kurt$), and the drawdown (DD) of volatility managed option trading strategies, as well as their differences to the unmanaged strategies. The data covers 332 months per strategy during the period from January 1990 to August 2017. Standard errors of $\Delta skew$, $\Delta kurt$, and ΔDD are obtained using a circular block bootstrap with 10'000 repetitions and are reported in parentheses.

	$skew$	$kurt$	DD	$\Delta skew$		$\Delta kurt$		ΔDD	
Cov. Call ATM (BXM)	-0.97	1.87	29.18	0.31	(0.47)	-3.21*	(1.65)	-7.76	(7.17)
Cov. Call 2% OTM (BXY)	-0.58	0.62	36.07	0.31	(0.31)	-2.29**	(0.98)	-5.42	(4.52)
Cov. Call 30Δ (BXMD)	-0.53	0.79	35.34	0.30	(0.31)	-1.85**	(0.88)	-8.43	(6.98)
Put-Write (PUT)	-1.50	4.78	29.24	0.41	(0.74)	-4.59	(3.66)	-3.82	(3.83)
Covered Combo (CMBO)	-0.89	1.56	31.80	0.38	(0.42)	-3.14**	(1.56)	-7.62	(6.92)
Collar (CLL)	-0.02	-0.06	31.29	0.15	(0.11)	0.15	(0.29)	-6.65	(6.87)
Zero-Cost Collar (CLLZ)	-0.44	0.67	38.43	0.49	(0.37)	-2.21*	(1.30)	-12.58	(12.2)
Protective Put (PPUT)	-0.15	0.06	36.28	0.18	(0.14)	-0.27	(0.29)	-13.03	(15.8)
Butterfly (BFLY)	-0.12	-0.16	35.03	-0.18*	(0.09)	0.21	(0.21)	1.03	(19.47)
Condor (CNDR)	-1.84	5.95	15.53	0.26	(0.24)	0.08	(0.59)	-1.07	(1.96)
S&P500 Index	-0.35	0.32	44.92	0.24	(0.24)	-1.04	(0.66)	-16.88	(15.73)

*Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$*

they are weaker. For example, the PPUT strategy benefits from an increase in skewness by 0.15, a reduction in kurtosis of 0.27, and a reduction in maximum drawdown by 13.03 percentage points. Moreover, both beta and downside beta are decreased by 0.05 and 0.20 respectively. These findings are in line with the previous argumentation that volatility control is less effective with strategies that have limited downside risk.

Finally, volatility control is less effective in reducing the risk of butterfly and condor strategies. The risk of the butterfly strategy remains practically unchanged and the absolute skewness and kurtosis even increase by 0.18 and 0.21 respectively. Volatility control is more effective for the condor strategy, leading to a slight reduction in absolute skewness and kurtosis. The drawdown and downside beta of the BFLY strategy remains unchanged while the CNDR strategy benefits from a reduction of downside beta of 0.23.

3.5 What drives the success of volatility control?

The results in section 3.3 show that volatility control generates alpha and increases the risk-return trade-off of option trading strategies. In this section I analyze the drivers of the success of volatility control. In particular, I first examine the impact of return characteristics on the success of volatility control as described by Perchet et al. (2016). Furthermore, I decompose option returns into a static equity, a dynamic equity, and a volatility component as suggested by Israelov & Nielsen (2015a).

3.5.1 The impact of return characteristics

Perchet et al. (2016) links volatility control to three phenomena which are commonly observed in financial markets. First, the volatility of returns has to cluster, i.e. it has to have a positive autocorrelation. Volatility control increases Sharpe ratios because the exposure to a strategy is reduced during times of high volatility thus reducing the overall volatility of the strategy. Therefore, volatility control in the

presence of volatility clustering leads to a higher Sharpe ratio by reducing volatility. If volatility is not autocorrelated, last period's volatility is a poor predictor of the next period's volatility and no effective volatility reduction takes place. However, investors might still be successful at volatility control by using a more elaborate volatility forecasting model. Second, Perchet et al. (2016) link volatility control to the negative autocorrelation between returns and volatility. This negative correlation between returns and volatility is often referred to as leverage effect. If a financial time series exhibits a leverage effect, volatility control helps to avoid negative returns during times of high volatility. Therefore, volatility control in the presence of a leverage effect leads to a higher Sharpe ratio by reducing negative returns. Third, fat tails in the return distribution increase the efficiency of volatility control. Perchet et al. (2016) argue that when fat tails are present, reducing the exposure to a strategy in times of high volatility especially help to reduce drawdowns.

Figure 4 shows the relationship between the alpha of the volatility controlled option strategy benchmarks over their unmanageed counterparts in relation to the previously described effects. Subfigure a) shows the the relation between the autocorrelation of volatility in the unmanaged strategies[9] The results show a clear dependency between autocorrelation of volatility and alpha. Strategies with high autocorrelation of volatility in their unmanaged time-series have significantly higher alphas than series with low autocorrelation (BFLY and CNDR). However, among strategies with high autocorrelation in their volatilities the relationship between volatility clustering and alpha breaks down which could indicate that autocorrelation is a necessary but not a sufficient condition. As soon as the necessary level of autocorrelation is reached, the other two effects become more important. Figure b) shows a strong relation between alpha and the leverage effect. Option series with lower ("more negative") correlations between returns and volatility have higher alphas. Figure c) and d) examine the re-

[9]Note that the displayed autocorrelations of volatility are relatively high by construction because volatilities are computed as an exponentially weighted moving average with the center of mass at 20 days as in Moskowitz et al. (2012).

lation between fat tails and the efficiency of volatility control in option strategies. Perchet et al. (2016) simulate fat tails by drawing residuals in a GARCH model from a t-distribution which induces both skewness and kurtosis. The results show that strategies with lower skewness and higher kurtosis have higher alphas. In conclusion the results are in line with the findings of Perchet et al. (2016).

3.5.2 Decomposing option returns

The results in section 3.3 show that volatility control generates alpha and increases the risk-return trade-off of option trading strategies. However, it is not directly clear if volatility control generates alpha by controlling the exposure to the equity risk premium of the underlying S&P500 or by controlling the exposure to the variance risk premium. The returns of the Cboe option strategy benchmarks are driven by price changes in the S&P500 and by changes in volatility. Thus, the question arises whether the success of volatility control is related to controlling volatility in the S&P500 or by controlling volatility of volatility arises. To ensure that the previous results are not driven by simply controlling the S&P500, I use the decomposition of Israelov & Nielsen (2015a) who provide an elegant way to decompose option returns. Their main insight is that option trading strategies are composed of three different return processes. The first two are the return generated by changes in the underlying and the return generated by collecting or paying the variance risk premium of the underlying. The third component is the dynamic exposure to the equity risk premium. Israelov & Nielsen (2015a) argue that the exposure to the underlying is not constant for option trading strategies. As options move into- or out-of-the money, their delta, and thus their exposure to the underlying, changes. For example, the equity exposure of a covered call strategy is dependent on the price change of the underlying after the call options were written. If the value of the underlying increases, the call moves into-the-money and the delta of the call option increases towards 1. Thus, the covered call strategy delta approaches 0, because it consists of an investment in the underlying

Figure 4: Return characteristics and volatility control

This figure shows the alphas of the volatility controlled option strategies over their uncontrolled benchmarks in relation to return characteristics of the underlying return series for all 10 Cboe option strategy benchmarks and the S&P500. In particular, Figure a) shows the the relation between volatility clustering and alpha, b) between the leverage effect and alpha, c) between skewness and alpha, and d) between excess kurtosis and alpha. Alphas are annualized and reported in percent.

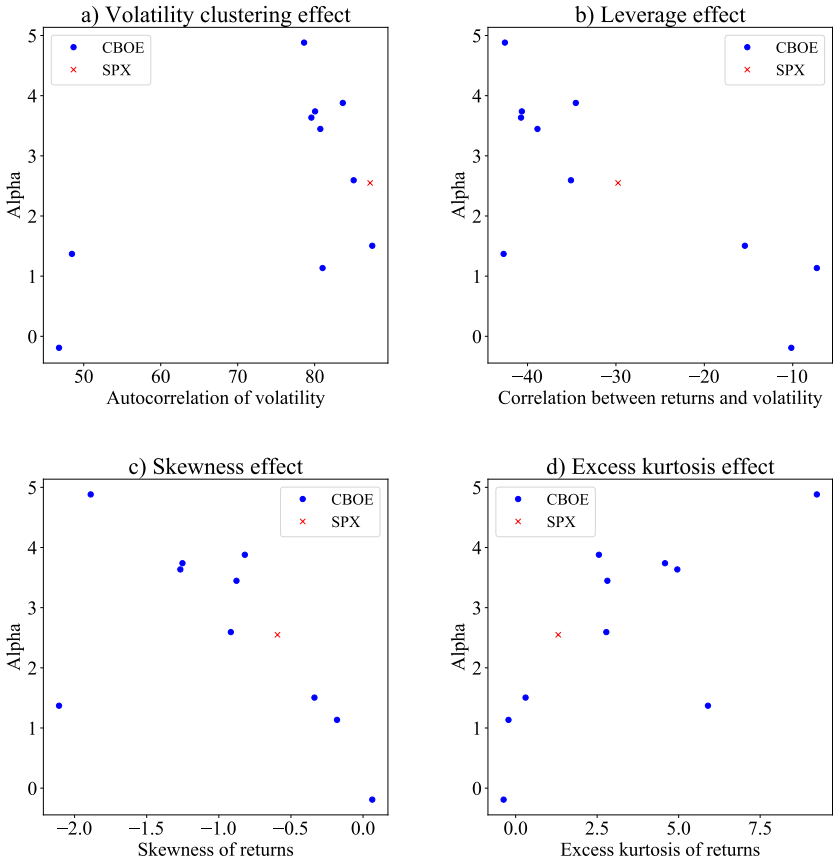

(delta 1) and a short call (approaching delta 1) which cancel each other out as the call is a short position. In contrast, if the value of the underlying decreases, the delta of the options approaches 0 and thus the portfolio delta approaches 1. Israelov & Nielsen (2015a) show that this dynamic equity component carries most of the risk and little return in covered call strategies. I generalize their approach as the sample contains also other strategies than covered calls. Note that the following equations show the decomposition on a total return basis instead of an excess return basis for clarity. The full methodology for the decomposition of any option strategy is presented in the Appendix. The returns of the Cboe option strategies are decomposed as follows:

$$ret_{tot,t} = \frac{port_t - port_{t-1}}{port_{t-1}} \tag{30}$$

$$ret_{eq,t} = \bar{\Delta}_{adj} \left(\frac{spx_t + div_t - spx_{t-1}}{spx_{t-1}} \right) \tag{31}$$

$$ret_{dyn,t} = (\Delta_{adj,t-1} - \bar{\Delta}_{adj}) \left(\frac{spx_t + div_t - spx_{t-1}}{spx_{t-1}} \right) \tag{32}$$

$$ret_{opt,t} = ret_{tot,t} - \Delta_{adj,t-1} \left(\frac{spx_t + div_t - spx_{t-1}}{spx_{t-1}} \right) \tag{33}$$

$$\tag{34}$$

where Δ refers to the portfolio's properly levered delta defined as:

$$\Delta_{adj,t} = \Delta_t \frac{spx_t}{port_t} \tag{35}$$

and Δ_t is the unlevered delta of the strategy computed as the sum of all portfolio component deltas multiplied by +/- 1 depending on whether the position is long or short. For example, a covered call strategy has an unlevered delta of $1 - \Delta_{Call}$ as it consists of long position in the S&P500 (delta 1) and a short position of an ATM call (delta $-\Delta_{Call}$). The total return (ret_{tot}) of an option strategy is defined as the percentage change in the strategy's portfolio ($port$) containing the position(s) in the option(s) and potential collaterals, e.g. a position in the S&P500 in case of covered call indices or the money market position in case of a put-write strategy. The equity component

of the return (ret_{eq}) is defined as the average delevered portfolio delta of the strategy ($\bar{\Delta}_{adj}$) multiplied by the return of the S&P500 plus any dividend payments. spx_t is the current value of the S&P500 and div_t are the dividends paid expressed in index points. Therefore, the equity return component of an option strategy is the part of the return that is attributable to the static exposure to the market equity risk premium. In contrast, ret_{dyn} is the dynamic exposure to the equity risk premium caused by the time-varying delta of the option strategies. The third component of the return is the delta hedged option strategy return ($ret_{opt,t}$). This return is constructed to resemble the return a trader would generate by statically delta hedging a portfolio of options at the end of each trading day. Delta hedged options are often considered to be investments into volatility as they are relatively unaffected by changes in the underlying and mostly capture changes in volatility (e.g. Goyal & Saretto, 2009).

Cboe does not deliver the delta of the option trading benchmarks. Therefore, I replicate the trading strategies according to the methodologies published on the Cboe website (see Appendix for details). They are decomposed into their equity, dynamic equity and volatility component on a daily basis, and are then aggregated to obtain monthly returns. Therefore, I obtain four return series for each option strategy on a monthly frequency: the returns of the three components and the total return of the option strategy. I then apply the volatility control methodology from Section 3.3 to each individual component of the strategy. Specifically, I use the realized volatility of the total (undecomposed) returns to control the exposure to all three components of each option trading strategy. As in section 3.3, the constant c is set for each option strategy so that the volatility controlled total return series has the same volatility as its uncontrolled counterpart. The individual components of each strategy are then scaled with the same value of c to make the alphas comparable in magnitude. Afterwards, the success of volatility control of the individual components is evaluated by computing the alpha of the volatility controlled strategy components over the uncontrolled components.

The results are shown in Table 29. Panel A shows the alphas of the volatility controlled individual components over their uncontrolled counterparts[10]. Panel B shows the betas of the volatility controlled strategies with respect to their uncontrolled counterparts. Evidence of volatility control is present for the replicated strategies during the sample period from January 1996 to August 2017. The alphas for the total return of put- and buy-write strategies range between 3.04 for the BXY strategy to 3.98 for the PUT strategy. The decomposition reveals that the alpha of volatility control is driven by both controlling the exposure to the variance risk premium and the exposure to the equity risk premium. For example, controlling the volatility component of the PUT strategy delivers a significant alpha of 1.08% per year, and controlling the equity component delivers a significant alpha of 1.61% per year. The alpha generated by controlling the equity component is less statistically significant, but larger in magnitude and thus economically more significant. This finding holds for all put- and buy-write strategies. Volatility control in put- and buy-write strategies appears to control exposures to both equity and volatility risk premia.

The results of other replicated strategies matches the previous results. All other replicated strategies generate positive but insignificant alphas. Alphas for defensive strategies, such as the PPUT strategy, are smaller and they generate a negative albeit insignificant negative alpha from controlling the option component. A potential explanation for the poor efficacy of volatility control of the PPUT AND CLL strategies is that these strategies have a negative exposure to the variance risk premium by construction. Therefore, the alpha generated by volatility control is negative which partially offsets the positive alpha generated by controlling the exposure to the equity risk premium. In conclusion, the results show that volatility control in option trading is statistically more successful in controlling the exposure to the the volatility risk premium but economically more successful by controlling the exposure to the equity

[10]Note that the alphas of the equity components are different for each strategy because each of them is timed with the volatility of the undecomposed strategy returns.

Table 29: The drivers of volatility control

This table shows the alpha of volatility controlled components of option trading strategies. ret_{tot} is the total return of the strategy consisting of the sum of all three sub-parts, ret_{eq} is the return attributable to the equity risk premium, ret_{dyn} is the return attributable to the dynamic exposure of the equity risk premium, and ret_{vol} is the return attributable to the variance risk premium. All components are managed with the realized volatility of the undecomposed total return. Standard errors are Newey & West (1987) adjusted with 5 lags and are reported in parentheses. Alphas are annualized and reported in percent. The data covers the period from January 1996 to August 2017.

	Panel A: Alphas				**Panel B: Betas**			
	ret_{tot}	ret_{eq}	ret_{dyn}	ret_{opt}	ret_{tot}	ret_{eq}	ret_{dyn}	ret_{opt}
Cov. Call ATM (BXM)	3.22^{**}	1.91^{*}	0.20	0.64^{**}	0.86^{***}	0.97^{***}	0.93^{***}	1.09^{***}
	(1.61)	(1.0)	(0.66)	(0.32)	(0.0)	(0.0)	(0.11)	(0.0)
Cov. Call 2% OTM (BXY)	3.04^{*}	2.27^{*}	−0.25	0.81^{***}	0.89^{***}	0.93^{***}	0.92^{***}	0.95^{***}
	(1.65)	(1.16)	(−0.46)	(0.23)	(0.0)	(0.0)	(0.0)	(0.0)
Cov. Call 30Δ (BXMD)	3.42^{*}	2.55^{**}	−0.24	0.62^{**}	0.88^{***}	0.94^{***}	0.88^{***}	1.08^{***}
	(1.75)	(1.29)	(−0.51)	(0.25)	(0.0)	(0.0)	(0.14)	(0.0)
Put-Write (PUT)	3.98^{**}	1.61^{*}	0.37	1.08^{**}	0.85^{***}	0.99^{***}	0.94^{***}	1.00^{***}
	(1.89)	(0.95)	(0.69)	(0.53)	(0.12)	(0.0)	(0.12)	(0.12)
Cov. Combo (CMBO)	3.16^{**}	1.61^{*}	0.05	1.00^{***}	0.87^{***}	0.95^{***}	0.91^{***}	0.94^{***}
	(1.59)	(0.97)	(0.54)	(0.24)	(0.00)	(0.00)	(0.00)	(0.00)
Collar (CLL)	0.81	1.09	−0.06	−0.09	0.86^{***}	0.81^{***}	0.75^{***}	0.76^{***}
	(0.56)	(0.79)	(−0.24)	(−0.13)	(0.00)	(0.00)	(0.00)	(0.00)
ZC Collar (CLLZ)	2.32	1.85	−0.24	0.46^{***}	0.86^{***}	0.89^{***}	0.80^{***}	0.95^{***}
	(1.45)	(1.16)	(−0.31)	(0.13)	(0.00)	(0.00)	(0.15)	(0.00)
Prot. Put (PPUT)	1.31	1.49	−0.07	−0.17	0.99^{***}	0.93^{***}	0.86^{***}	0.86^{***}
	(0.83)	(1.09)	(−0.36)	(−0.16)	(0.00)	(0.00)	(0.11)	(0.00)
Butterfly (BFLY)	0.04	−0.12	0.12	−0.04	0.97^{***}	0.89^{***}	0.94^{***}	0.97^{***}
	(0.54)	(−0.2)	(0.6)	(−0.39)	(0.0)	(0.0)	(0.0)	(0.0)
Condor (CNDR)	1.43	−0.04	0.39	1.01^{*}	0.91^{***}	0.91^{***}	0.89^{***}	0.96^{***}
	(0.88)	(−0.05)	(0.6)	(0.59)	(0.0)	(0.0)	(0.0)	(0.0)
S&P500 Index	2.17	2.17	–	–	0.91^{***}	0.91^{***}	–	–
	(1.57)	(1.57)	–	–	(0.00)	(0.00)	–	–

Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$

risk premium. Therefore, volatility control of option strategies seems to generate alpha by adjusting the exposure to the variance risk premium and the equity risk premium.

3.6 The impact of transaction costs

Volatility control requires that portfolio positions are rebalanced once per month and thus induces additional transaction costs to a trading strategy. Transaction costs

may reduce the alpha of any strategy significantly. In Table 30, the effect of transaction costs on the profitability of volatility control is shown. Transaction costs are applied as a percent of the change in the position size. Therefore, the results are independent of the invested amount or the choice of the constant c. In addition to the alphas of volatility controlled option trading strategies over the unmanaged counterparts, Table 30 also shows the transaction costs at which the alpha of the volatility controlled option strategies is zero.

The findings show that strategies with higher alphas can also sustain higher transaction costs. For example, put- and buy-write strategies can still generate significant alphas at transaction costs of 0.25% per transaction, and can sustain maximum transaction costs between 0.84% (BXM) and 1.15% (BXMD). Protective put (PPUT) and collar (CLL) strategies can sustain transaction costs of 0.81% and 0.61% respectively. Moreover, a condor strategy can sustain maximum transaction costs of 0.30%.

The results are conservative estimates as the transaction costs can be significantly reduced by rebalancing the portfolios on option expiration days. Option trading strategies have to be rolled at the expiration date of options, typically on the third Friday of each month. If investors choose to rebalance on expiration days, a majority of the additional rebalancing costs induced by volatility control can be saved as a new set of options has to be bought in any case.

3.7 Conclusion

In this study, I show that option trading strategies can be managed by reducing exposure when volatility is high and increasing exposure if volatility is low. This volatility control generates sizable alphas for 9 out of 10 Cboe option strategy benchmarks. Volatility control is particularly successful for put- and buy-write strategies which usually carry high levels of downside risk. Volatility controlled versions of these strategies generate significant alphas over the uncontrolled series and have increased

Table 30: Impact of transaction costs

This table shows the impact of transaction costs on volatility managed option trading strategies. The columns show the alphas of the volatility managed strategies over their unmanaged counterparts for different levels of transaction costs. The last column shows the transaction costs in percent at which the managed strategies have an alpha of 0. The data covers 332 months per strategy during the period from January 1990 to August 2017. Standard errors of alphas are Newey & West (1987) adjusted with 5 lags and are not reported to keep the table concise. Alphas are annualized and reported in percent.

	α at transaction costs of				
	0.0%	0.1%	0.25%	0.5%	$maxTC$
Covered Call ATM (BXM)	3.64^{**}	3.20^{**}	2.55^{*}	1.47	0.84
Covered Call 2% OTM (BXY)	3.45^{**}	3.14^{**}	2.67^{*}	1.90	1.12
Covered Call 30Δ (BXMD)	3.88^{**}	3.54^{**}	3.03^{*}	2.18	1.15
Put-Write (PUT)	4.88^{***}	4.34^{**}	3.53^{**}	2.18	0.91
Covered Combo (CMBO)	3.74^{**}	3.36^{**}	2.78^{*}	1.83	0.98
Collar (CLL)	1.13^{*}	0.95	0.67	0.21	0.61
Zero-Cost Collar (CLLZ)	2.59^{*}	2.32^{*}	1.91	1.23	0.95
Protective Put (PPUT)	1.51^{*}	1.32^{*}	1.04	0.58	0.81
Butterfly (BFLY)	-0.19	-0.45	-0.83^{*}	-1.47^{***}	0.00
Condor (CNDR)	1.37^{**}	0.93	0.26	-0.84	0.30
S&P500 Index	2.55^{*}	2.30	1.94	1.33	1.04

Stars indicate statistical significance: * - $p \leq 0.1$, ** - $p \leq 0.05$, *** - $p \leq 0.01$

Sharpe ratios. Moreover, they have significantly lower downside risk, lower drawdowns and more normal returns.

Option trading strategies are profitable because they collect the equity- and the variance risk premium. I decompose the returns of the option strategies in the sample and find that volatility control generates alpha by controlling the exposure to the variance risk premium. Therefore, the effect is distinct from simply controlling exposure to the S&P500. These results are robust to transaction costs and show that option investors can significantly improve their performance by applying volatility control to option strategies.

3.8 Appendix: Replication and Decomposition of Cboe Strategies

This study aims to decompose the Cboe strategy returns. Cboe provides returns of different option-strategies, but does not provide data on the delta of the portfolio which is necessary for the decomposition. Therefore, I replicate the 10 Cboe option strategy benchmarks as closely as possible. Cboe publishes methodological whitepapers on their website which describe the methodology Cboe uses to compute the strategy returns.

The replication strategy is as follows: On each roll day, I select a subsample of options which a) expire on the next expiration day, b) have positive open interest and volume, and c) have data available on at least 90% of all trading days until the next expiration. If no option matches the targeted portfolio position exactly the first option which is further out-of-the-money, i.e. above (below) the targeted price for calls (puts), is selected. For example, the BXM put write strategy requires to sell an at-the-money call and thus the call option with the first available strike above the current S&P500 price is selected. This selection procedure applies to the BXM, BXY, PUT, CMBO, CLL, CLLZ, PPUT, and BFLY indices as outlined by Cboe. The only exceptions are the BXMD and CNDR indices which select the closest option to the targeted strike from options that can be further in- or out-of-the-money. The option is then bought (sold) at the ask (bid) price and is held in the portfolio at the mid price. The option is then held until the last day before expiration, typically the third Thursday every month, and the position is closed by selling (buying) the option at the last bid (ask) price. Simultaneously, new option positions are entered and the procedure is repeated. This creates a minor difference to the Cboe methodology which rolls options at at 11am on the expiration date. Therefore, options are always rolled on the last day before expiration, except for the first day in the sample, January 04, 1996, where portfolio positions are entered for the first time and are held until the next expiration day. Moreover, missing values of delta and prices are forward filled, i.e. if the delta on a certain day is missing, it is assumed that it still equals the last

valid observation. However, S&P500 options are generally liquid and forward filling does not change any of the results.

In addition, on each roll day the collateral of the strategy is updated. Cboe option benchmarks are fully collaterized. Strategies selling calls, e.g. buy-write strategies, invest into the S&P500 to cover the liability from the call option and strategies selling puts invest an amount equal to the strike price of the option into a money-market portfolio, earning the 3 month T-bill rate. An exception are the BFLY and CNDR strategy which invest 10 times their maximum liability into a money-market account. The maximum liability of BFLY and CNDR options is calculated as the maximum of the difference between the strikes of the puts and the difference between the strikes of the calls.

The above replication methodology yields a collection of time-series, each spanning the period from buying the option to the last trading day of the option. On the day before expiration, the entire portfolio is liquidated and positions in a the new portfolio are entered. Liquidating the entire portfolio on the day before the expiration day is necessary due to data limitations and also avoids a special treatment of returns on the expiration day. The total value of the portfolio is calculated at the end of each trading day as:

$$port_t = opt_t + eq_t + cash_t \tag{36}$$

$$\tag{37}$$

where opt_t is the value of the current option position. opt_t is negative for short positions, i.e. if a put with a price of 5 dollar is sold opt_t would take a value of -5. Thus if the price of the option declines, the position increases in value. eq_t is the value of the current equity holdings. For example in the replicated BXM strategy eq_t would equal the closing price of the S&P500. In contrast, for the PUT strategy eq_t is always 0 as the PUT strategy does not hold a position in the S&P500. $cash_t$ is the

value of the current cash position, i.e. the value of the collateral and any proceeds from selling options. The cash position earns interest at the 3 month t-bill rate. For the protective put PPUT strategy this position is 0 as it does not sell options and also does not require a collateral. The total excess return of the strategy is thus:

$$r_{tot,t} = \frac{opt_t + eq_t + div_{port,t} + (1 + r_{cash,t})cash_t}{opt_{t-1} + eq_{t-1} + cash_{t-1}} - 1 - r_{cash,t} \tag{38}$$

where div_t are all dividends earned during day t. If the portfolio does not contain an equity position, e.g. in case of the PUT index, $div_{port,t}$ is 0. Once time-series of portfolio values, returns and deltas are obtained, the decomposition of returns is relatively straight-forward.

$$ret_{eq,t} = \bar{\Delta}_{adj}\left(\frac{spx_t + div_{spx,t}}{spx_{t-1}} - r_{cash,t}\right) \tag{39}$$

$$= \bar{\Delta}\frac{spx_t + div_{spx,t} - spx_{t-1}}{port_{t-1}} - \bar{\Delta}_{adj} r_{cash,t} \tag{40}$$

$$ret_{dyn,t} = (\Delta_{adj,t-1} - \bar{\Delta}_{adj})\left(\frac{spx_t + div_{spx,t} - spx_{t-1}}{spx_{t-1}} - r_{cash,t}\right) \tag{41}$$

$$= (\Delta_{t-1} - \bar{\Delta})\frac{spx_t + div_{spx,t} - spx_{t-1}}{port_{t-1}} - (\Delta_{adj,t-1} - \bar{\Delta}_{adj})r_{cash,t} \tag{42}$$

where $ret_{eq,t}$ and $ret_{dyn,t}$ are the dynamic and static equity part of the return respectively. The adjusted Δ is the properly levered delta:

$$\Delta_{adj,t} = \frac{spx_t}{port_t}\Delta_t \tag{43}$$

where Δ_t is the aggregated portfolio delta, i.e. the sum of all portfolio components. $\bar{\Delta}$ is the sample mean of deltas. The return attributable to the option portion of the portfolio can then simply be defined as the residual of subtracting the equity components from the total return:

$$ret_{opt,t} = ret_{tot,t} - ret_{eq,t} - ret_{dyn,t} \tag{44}$$

$$= \frac{(opt_t - opt_{t-1}) + (eq_t + div_{port,t} - eq_{t-1}) + (1 + r_{cash,t})cash_t}{opt_{t-1} + eq_{t-1} + cash_{t-1}} \tag{45}$$

$$- \frac{\Delta_{t-1}(spx_t + div_{spx,t} - spx_{t-1})}{opt_{t-1} + eq_{t-1} + cash_{t-1}} - (1 - \Delta_{adj,t-1})r_{cash,t}$$

References

Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, *53*(2), 499–547.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, *5*(1), 31–56.

An, B.-J., Ang, A., Bali, T. G., & Cakici, N. (2014). The joint cross section of stocks and options. *Journal of Finance*, *69*(5), 2279–2337.

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance*, *61*(1), 259–299.

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2009). High idiosyncratic volatility and low returns: International and further us evidence. *Journal of Financial Economics*, *91*(1), 1–23.

Augustin, P., Brenner, M., & Subrahmanyam, M. (2015). Informed options trading prior to m&a announcements: Insider trading? *Working Paper*.

Bakshi, G., & Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies*, *16*(2), 527–566.

Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, *16*(1), 101–143.

Bakshi, G., & Madan, D. (2000). Spanning and derivative-security valuation. *Journal of Financial Economics*, *55*(2), 205–238.

Balanda, K. P., & MacGillivray, H. (1988). Kurtosis: a critical review. *The American Statistician*, *42*(2), 111–119.

Bali, T. G., & Murray, S. (2013). Does risk-neutral skewness predict the cross-section of equity option portfolio returns? *Journal of Financial and Quantitative Analysis*, *48*(4), 1145–1171.

Barberis, N., & Huang, M. (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, *98*(5), 2066–2100.

Barone Adesi, G. (2016). Var and cvar implied in option prices. *Journal of Risk and Financial Management*, *9*(1), 2.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies*, *9*(1), 69–107.

Birru, J., & Figlewski, S. (2012). Anatomy of a meltdown: The risk neutral density for the s&p 500 in the fall of 2008. *Journal of Financial Markets*, *15*(2), 151–180.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Bliss, R. R., & Panigirtzoglou, N. (2002). Testing the stability of implied probability density functions. *Journal of Banking & Finance*, *26*(2), 381–422.

Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, *51*(4), 621–651.

Buss, A., & Vilkov, G. (2012). Measuring equity risk with option-implied correlations. *Review of Financial Studies*, *25*(10), 3113–3140.

Cao, M., & Wei, J. (2010). Option market liquidity: Commonality and other characteristics. *Journal of Financial Markets*, *13*(1), 20–48.

Carr, P., & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, *22*(3), 1311–1341.

Carr, P., & Wu, L. (2017). Leverage effect, volatility feedback, and self-exciting market disruptions. *Journal of Financial and Quantitative Analysis*, 1–38.

Chakravarty, S., Gulen, H., & Mayhew, S. (2004). Informed trading in stock and option markets. *Journal of Finance*, *59*(3), 1235–1257.

Chang, B.-Y., Christoffersen, P., Jacobs, K., & Vainberg, G. (2011). Option-implied measures of equity risk. *Review of Finance*, *16*(2), 385–428.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, *74*(368), 829–836.

Conrad, J., Dittmar, R. F., & Ghysels, E. (2013). Ex ante skewness and expected stock returns. *Journal of Finance*, *68*(1), 85–124.

Coval, J. D., & Shumway, T. (2001). Expected option returns. *Journal of Finance*, *56*(3), 983–1009.

Cremers, M., & Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, *45*(2), 335–367.

Daniel, K., & Titman, S. (2006). Market reactions to tangible and intangible information. *Journal of Finance*, *61*(4), 1605–1643.

Dennis, P., & Mayhew, S. (2009). Microstructural biases in empirical tests of option pricing models. *Review of Derivatives Research*, *12*(3), 169–191.

Easley, D., O'Hara, M., & Srinivas, P. S. (1998). Option volume and stock prices: Evidence on where informed traders trade. *Journal of Finance*, *53*(2), 431–465.

Eraker, B. (2013). The performance of model based option trading strategies. *Review of Derivatives Research*, *16*(1), 1–23.

Fama, E. F., & French, K. R. (2016). Dissecting anomalies with a five-factor model. *Review of Financial Studies*, *29*(1), 69–103.

Fan, J. (1993). Local linear regression smoothers and their minimax efficiencies. *Annals of Statistics*, *21*(1), 196–216.

Fleming, J., Kirby, C., & Ostdiek, B. (2001). The economic value of volatility timing. *Journal of Finance*, *56*(1), 329–352.

Fleming, J., Kirby, C., & Ostdiek, B. (2003). The economic value of volatility timing using realized volatility. *Journal of Financial Economics*, *67*(3), 473–509.

Frazzini, A., & Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, *111*(1), 1–25.

Gao, G. P., Gao, P., & Song, Z. (2018). Do hedge funds exploit rare disaster concerns? *The Review of Financial Studies*, *31*(7), 2650–2692.

Gao, G. P., Lu, X., & Song, Z. (2018). Tail risk concerns everywhere. *forthcoming in Management Science*.

Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. *Review of Financial Studies*, *22*(10), 4259–4299.

Gatheral, J. (2011). *The volatility surface: A practitioner's guide*. John Wiley & Sons.

Gkionis, K., Kostakis, A., Skiadopoulos, G. S., & Stilger, P. S. (2018). Positive stock information in out-of-the-money option prices. *Working Paper*.

Goncalves-Pinto, L., Grundy, B. D., Hameed, A., van der Heijden, T., & Zhu, Y. (2016). Why do option prices predict stock returns? the role of price pressure in the stock market. *Working Paper*.

Goyal, A., & Saretto, A. (2009). Cross-section of option returns and volatility. *Journal of Financial Economics*, *94*(2), 310–326.

Griffin, J. M., & Shams, A. (2017). Manipulation in the vix? *The Review of Financial Studies*, *31*(4), 1377–1417.

Groeneveld, R. A., & Meeden, G. (1984). Measuring skewness and kurtosis. *The Statistician*, *33*(4), 391–399.

Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, *55*(3), 1263–1295.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, *6*(2), 327–343.

Hinkley, D. V. (1975). On power transformations to symmetry. *Biometrika*, *62*(1), 101–111.

Israelov, R. (2017). Putwrite versus buywrite: Yes, put-call parity holds here too. *AQR Working Paper*.

Israelov, R., & Nielsen, L. N. (2014). Covered call strategies: One fact and eight myths. *Financial Analysts Journal*, *70*(6), 23–31.

Israelov, R., & Nielsen, L. N. (2015a). Covered calls uncovered. *Financial Analysts Journal*, *71*(6), 44–57.

Israelov, R., & Nielsen, L. N. (2015b). Still not cheap: Portfolio protection in calm markets. *Journal of Portfolio Management*, *41*(4), 108–120.

Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance*, *51*(5), 1611–1631.

Jiang, G., Lee, C. M., & Zhang, Y. (2005). Information uncertainty and expected returns. *Review of Accounting Studies*, *10*(2), 185–221.

Jiang, G. J., & Tian, Y. S. (2007). Extracting model-free volatility from option prices: An examination of the vix index. *Journal of Derivatives*, *14*(3), 35–60.

Jin, W., Livnat, J., & Zhang, Y. (2012). Option prices leading equity prices: do option traders have an information advantage? *Journal of Accounting Research*, *50*(2), 401–432.

Jones, M., & Pewsey, A. (2009). Sinh-arcsinh distributions. *Biometrika*, *96*(4), 761–780.

Jurek, J. W. (2014). Crash-neutral currency carry trades. *Journal of Financial Economics*, *113*(3), 325–347.

Kozhan, R., Neuberger, A., & Schneider, P. (2013). The skew risk premium in the equity index market. *Review of Financial Studies*, *26*(9), 2174–2203.

Kraus, A., & Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *Journal of Finance*, *31*(4), 1085–1100.

Lettau, M., Maggiori, M., & Weber, M. (2014). Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, *114*(2), 197–225.

Li, Q., & Racine, J. (2004). Cross-validated local linear nonparametric regression. *Statistica Sinica*, *14*(2), 485–512.

Lucca, D. O., & Moench, E. (2015). The pre-fomc announcement drift. *The Journal of Finance*, *70*(1), 329–371.

Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics*, *132*(1), 367–433.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, *3*(1), 125–144.

Mirkov, N., Pozdeev, I., & Söderlind, P. (2019). Verbal interventions and exchange rate policies: The case of swiss franc cap. *Journal of International Money and Finance*, *93*, 42–54.

Moreira, A., & Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance*, *72*(4), 1611–1644.

Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. *Journal of Financial Economics*, *104*(2), 228–250.

Neumann, M., & Skiadopoulos, G. (2013). Predictable dynamics in higher-order risk-neutral moments: Evidence from the s&p 500 options. *Journal of Financial and Quantitative Analysis*, *48*(03), 947–977.

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708.

Perchet, R., de Carvalho, R. L., Heckel, T., & Moulin, P. (2016). Predicting the success of volatility targeting strategies: Application to equities and other asset classes. *Journal of Alternative Investments*, *18*(3), 21–38.

Rehman, Z., & Vilkov, G. (2012). Risk-neutral skewness: Return predictability and its sources. *Working Paper*.

Ross, S. (2015). The recovery theorem. *Journal of Finance*, *70*(2), 615–648.

Ruppert, D. (1987). What is kurtosis? An influence function approach. *The American Statistician*, *41*(1), 1–5.

Schmid, F., & Trede, M. (2003). Simple tests for peakedness, fat tails and leptokurtosis based on quantiles. *Computational statistics & data analysis*, *43*(1), 1–12.

Schneider, P., Wagner, C., & Zechner, J. (2015). Low risk anomalies? *Working Paper*.

Sharpe, W. F. (1994). The sharpe ratio. *Journal of Portfolio Management*, *21*(1), 49–58.

Shimko, D. (1993). The bounds of probability. *RISK*, *6*, 33–37.

Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, *52*(1), 35–55.

Söderlind, P., & Svensson, L. (1997). New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics*, *40*(2), 383–429.

Song, Z., & Xiu, D. (2016). A tale of two option markets: Pricing kernels and volatility risk. *Journal of Econometrics*, *190*(1), 176–196.

Stilger, P. S., Kostakis, A., & Poon, S.-H. (2016). What does risk-neutral skewness tell us about future stock returns? *Management Science*, *Articles in Advance*.

Stone, C. J. (1977). Consistent nonparametric regression. *Annals of Statistics*, *5*(4), 595–620.

Tian, Y. S. (2011). Extracting risk-neutral density and its moments from american option prices. *Journal of Derivatives*, *18*(3), 17–34.

Ungar, J., & Moran, M. T. (2009). The cash-secured putwrite strategy and performance of related benchmark indexes. *The Journal of Alternative Investments*, *11*(4), 43–56.

Whaley, R. E. (2002). Return and risk of cboe buy write monthly index. *Journal of Derivatives*, *10*(2), 35–42.

Xing, Y., Zhang, X., & Zhao, R. (2010). What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, *45*(3), 641–662.

Zhang, X. (2006). Information uncertainty and stock returns. *Journal of Finance*, *61*(1), 105–137.

Curriculum Vitae

Alexander Feser

Education

2020: PhD in Finance, University of St. Gallen

2014: MA in Quantitative Economics and Finance, University of St. Gallen

2012: BA in Economics, University of St. Gallen

Research papers

Option-Implied Value-at-Risk and the Cross-Section of Stock Returns (joint with Manuel Ammann)

Robust Estimation of Risk-Neutral Moments (joint with Manuel Ammann)

Volatility Control of Option Strategies

Referee Activity

Journal of Banking and Finance

Financial Markets and Portfolio Management