

School of Finance



University of St.Gallen

ROBUST ESTIMATION OF RISK-NEUTRAL MOMENTS

MANUEL AMMANN

ALEXANDER FESER

WORKING PAPERS ON FINANCE NO. 2019/02

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

MARCH 20, 2019



Robust Estimation of Risk-Neutral Moments

MANUEL AMMANN and ALEXANDER FESER*

March 20, 2019

forthcoming 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.

*Manuel Ammann and Alexander Feser (alexander.feser@unisg.ch, +41 71 224 70 04) are at the Swiss Institute of Banking and Finance, University of St. Gallen, St. Gallen, Switzerland. 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.

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, Dittmar, & Ghysels, 2013; Kozhan, Neuberger, & Schneider, 2013; Jurek, 2014; Stiller, Kostakis, & Poon, 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, Kapadia, & Madan (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 estimators, 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. 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, Paul, & Söderlind (2018) 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

¹There 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).

often the last trade in each option contract happened at a different time of the day. This asynchronous 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 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 briefly reviews the risk-neutral moments of Bakshi et al. (2003) and introduces quantile moments. Section 3 discusses the different estimation techniques. Section 4 presents the methodology and the result of the horse-race between the different estimators. Section 5 presents our empirical results. Section 6 presents alternative specifications of the simulation study. Section 7 concludes.

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:

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

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

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.

2.1 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:

$$V[t, \tau] = \int_{S_t}^{\infty} \frac{2(1 - \ln[\frac{K}{S_t}])}{K^2} C(t, \tau, K) dK + \int_0^{S_t} \frac{2(1 - \ln[\frac{S_t}{K}])}{K^2} P(t, \tau, K) dK \quad (3)$$

$$W[t, \tau] = \int_{S_t}^{\infty} \frac{6\ln[\frac{K}{S_t}] - 3(\ln[\frac{K}{S_t}])^2}{K^2} C(t, \tau, K) dK - \int_0^{S_t} \frac{6\ln[\frac{S_t}{K}] - 3(\ln[\frac{S_t}{K}])^2}{K^2} P(t, \tau, K) dK \quad (4)$$

$$X[t, \tau] = \int_{S_t}^{\infty} \frac{12(\ln[\frac{K}{S_t}])^2 - 4(\ln[\frac{K}{S_t}])^3}{K^2} C(t, \tau, K) dK + \int_0^{S_t} \frac{12(\ln[\frac{S_t}{K}])^2 - 4(\ln[\frac{S_t}{K}])^3}{K^2} P(t, \tau, K) dK \quad (5)$$

The challenge in the estimation of central risk-neutral moments is to evaluate the integrals in Equations (3) - (5). 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:

$$Vol^{\mathbb{Q}} = \sqrt{e^{r\tau}V - \mu^2} \quad (6)$$

$$Skew^{\mathbb{Q}} = \frac{e^{r\tau}W - 3e^{r\tau}\mu V + 2\mu^3}{[e^{r\tau}V - \mu^2]^{3/2}} \quad (7)$$

$$Kurt^{\mathbb{Q}} = \frac{e^{r\tau}X - 4\mu W + 6e^{r\tau}\mu^2V - \mu^4}{[e^{r\tau}V - \mu^2]^2} \quad (8)$$

where

$$\mu = e^{r\tau} - 1 - e^{r\tau}\frac{V}{2} - e^{r\tau}\frac{W}{6} - e^{r\tau}\frac{X}{24} \quad (9)$$

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-quantile range is used to measure quantile volatility which is defined as:

$$QVol = Q(0.75) - Q(0.25) \quad (10)$$

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:

$$\text{QSkew}(p) = \frac{[Q(p) - Q(0.5)] - [Q(0.5) - Q(1 - p)]}{Q(p) - Q(1 - p)} \quad (11)$$

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:

$$\text{QKurt}(p) = \frac{Q(p) - Q(1 - p)}{Q(q) - Q(1 - q)} \quad (12)$$

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 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

²An 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).

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.

[Figure 1 approximately here.]

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 differences³.

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:

$$P(S, K, \sigma_K, \tau, r) = e^{-r\tau} K \Phi(-d_2(\sigma)) - S_0 \Phi(-d_1(\sigma)) \quad (13)$$

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:

$$F(S_\tau < K) = e^{r\tau} \Phi(-d_2(\sigma)) \quad (14)$$

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.

³The 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.

⁴Practitioners 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.

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:

$$SVIX_t = \frac{2e^{r\tau}}{\tau F_t(\tau)^2} \left(\int_0^{F_t(\tau)} P(t, \tau, K) dK + \int_{F_t(\tau)}^{\infty} C(t, \tau, K) dK \right) \quad (15)$$

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 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 3:

$$RIX = \frac{2e^{r\tau}}{\tau} \left(\int_0^{S_t} \frac{\ln(S_t/K)}{K^2} P(t, \tau, K) dK \right) \quad (16)$$

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:

$$VIX = \frac{2e^{r\tau}}{\tau} \left(\int_0^{S_t} \frac{1}{K^2} P(t, \tau, K) dK + \int_{S_t}^{\infty} \frac{1}{K^2} C(t, \tau, K) dK \right) \quad (17)$$

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.

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.

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 to 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, extrapolating 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. 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.

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.

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.

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:

$$\frac{dS_t}{S_t} = (r - \lambda\mu_J)dt + \sqrt{\sigma_t}dZ_t^1 + J_t dN_t \quad (18)$$

$$d\sigma_t = \kappa(\sigma_L - \sigma_t)dt + \nu\sqrt{\sigma_t}dZ_t^2 \quad (19)$$

$$\text{corr}(dZ^1, dZ^2) = \rho dt \quad (20)$$

$$\text{prob}(dN = 1) = \lambda dt \quad (21)$$

$$Z^1, Z^2 \sim N(0, 1) \quad (22)$$

$$\ln(1 + J) \sim N(\ln(1 + \mu_J) - 0.5\sigma_J^2, \sigma_J^2) \quad (23)$$

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

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 &

⁵The 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%.

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 than 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:

$$\tilde{P}_i = P_i(1 + \theta\eta) \quad \eta \sim N(0,1) \quad (24)$$

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 75th and 25th percentile. Quantile skewness is measured at the 10th and 90th percentile ($p = 0.9$) and quantile kurtosis is estimated with an outer range between the 95th and 5th percentile and an inner range between the 75th and 25th 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.

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 1 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 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 1 shows the results in tabular form for the domain half-widths of 10%, 50%, and 80%. In addition, Table 1 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 1 also contains estimates obtained from a local-constant kernel regression comparable in magnitude to the bias of the cubic smoothing spline with horizontal extrapolation.

[Figure 2 & Table 1 about here.]

The sensitivity of estimates to micro-structural noise is shown in Figure 3 and Table 2. 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 2). 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 2). 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

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

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, 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.

[Figure 3 and Table 2 about here.]

The approximation errors for different domain spacings (Table 3) 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.

[Table 3 about here.]

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, 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 4.

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.

[Table 4 about here.]

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.

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 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

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 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 5. The S&P 500 has a large

⁷Note 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, 2010). Naively joining put and call volatilities at-the-money will lead to a jump at-the-money.

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%.

5.2 Comparison of quantile and central moments

Summary statistics for option-implied moments and the VIX, RIX, and SVIX, are provided in Table 6. 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. Furthermore, 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.

[Table 6 about here.]

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 7 - 9 in the Appendix. The results mirror the previous results of the SVJ model. In Table 7 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 8 shows that methods with flat extrapolation are less affected by micro-structural noise. Errors from different strike price spacings (Table 9) 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 10 - 12 in the Appendix).

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 mo-

ments 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.

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.
- 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.
- 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.
- Carr, P., & Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3), 1311–1341.
- 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.
- Cremers, M., & Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2), 335–367.
- Dennis, P., & Mayhew, S. (2009). Microstructural biases in empirical tests of option pricing models. *Review of Derivatives Research*, 12(3), 169–191.
- Fan, J. (1993). Local linear regression smoothers and their minimax efficiencies. *Annals of Statistics*, 21(1), 196–216.
- 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. *Management Science*.
- Gatheral, J. (2011). *The volatility surface: A practitioner's guide*. John Wiley & Sons.
- Griffin, J. M., & Shams, A. (2017). Manipulation in the vix? *Working Paper*.
- Groeneveld, R. A., & Meeden, G. (1984). Measuring skewness and kurtosis. *The Statistician*, 33(4), 391–399.
- 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.
- Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. *Journal of Finance*, 51(5), 1611–1631.
- 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.


- 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.
- Li, Q., & Racine, J. (2004). Cross-validated local linear nonparametric regression. *Statistica Sinica*, 14(2), 485–512.
- 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., Paul, & Söderlind, P. (2018). Verbal interventions and exchange rate policies: The case of swiss franc cap. *Working Paper*.
- 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.
- 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.
- Shimko, D. (1993). The bounds of probability. *RISK*, 6, 33–37.

- 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*, 63(6), 1814–1834.
- Stone, C. J. (1977). Consistent nonparametric regression. *Annals of Statistics*, 5(4), 595–620.


Figures

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 75th and 25th percentile. Quantile skewness is defined by Hinkley's (1975) measure at the 90th percentile. Quantile kurtosis is defined by Ruppert's (1987) ratio of quantile ranges with an outer range between the 95th and 5th percentile and an inner range between the 75th and 25th 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


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




Panel B: Crisis Scenario - Different Inter- and Extrapolation




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


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



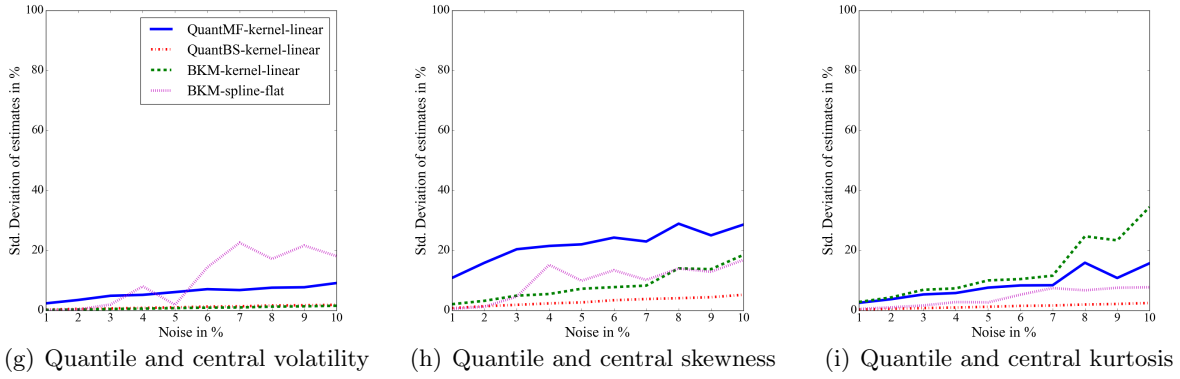
Panel B: Crisis Scenario - Different Inter- and Extrapolation



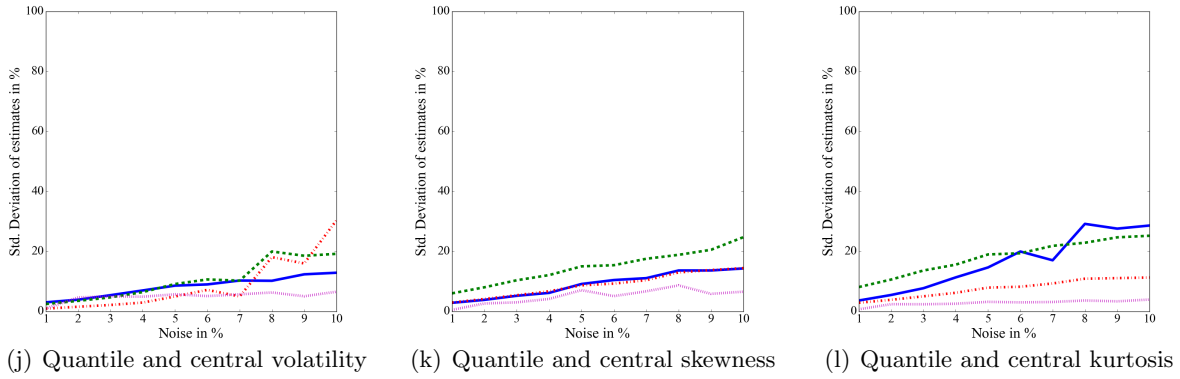
continued on the next page

Figure 3 continued

Panel C: Standard Scenario - Different Estimators



Panel D: Crisis Scenario - Different inter- and Extrapolation



Tables

Table 1: 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).

Domain half-width	Volatility			Skewness			Kurtosis		
	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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

Table 2: 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).

Noise	Volatility			Skewness			Kurtosis		
	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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

Table 3: 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).

Spacing	Volatility			Skewness			Kurtosis		
	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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

Table 4: 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									
Domain half-width	10%	50%	80%	10%	50%	80%	10%	50%	80%
<i>Standard scenario</i>									
kernel-linear	0.41	0.00	0.00	0.34	0.01	0.01	0.67	0.00	0.00
lkernel-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
<i>Crisis scenario</i>									
kernel-linear	0.41	0.00	0.00	0.34	0.01	0.01	0.67	0.00	0.00
lkernel-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
Panel B: Noise									
Noise	1%	5%	10%	1%	5%	10%	1%	5%	10%
<i>Standard scenario</i>									
kernel-linear	0.36	1.58	3.12	1.92	6.60	31.81	0.34	1.53	2.83
lkernel-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
<i>Crisis scenario</i>									
kernel-linear	5.90	19.20	70.19	78.89	287.17	1.72e3	1.55	4.94	10.75
lkernel-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									
Spacing	1%	2%	5%	1%	2%	5%	1%	2%	5%
<i>Standard scenario</i>									
kernel-linear	0.01	0.02	0.08	0.01	0.01	0.01	0.01	0.02	0.10
lkernel-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
<i>Crisis scenario</i>									
kernel-linear	0.00	0.01	0.01	0.10	0.08	0.06	0.01	0.00	0.02
lkernel-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

Table 5: 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.

	<i>Mean</i>	<i>Std.</i>	<i>Med.</i>	<i>Min</i>	<i>Max</i>
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

Table 6: 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 75th and 25th percentile. Quantile skewness is defined by Hinkley's (1975) measure at the 90th percentile. Quantile kurtosis is defined by Ruppert's (1987) ratio of quantile ranges with an outer range between the 95th and 5th percentile and an inner range between the 75th and 25th 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				
	<i>Mean</i>	<i>Std.</i>	<i>Med.</i>	<i>Min</i>	<i>Max</i>	<i>Mean</i>	<i>Std.</i>	<i>Med.</i>	<i>Min</i>	<i>Max</i>
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

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	

Appendix: Robustness tests

Table 7: 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).

Domain half-width	Volatility			Skewness			Kurtosis		
	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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 8: 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).

Noise	Volatility			Skewness			Kurtosis		
	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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	36051.02
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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	1554.00	1141.92
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 9: 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).

Spacing	Volatility			Skewness			Kurtosis		
	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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 10: 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).

Domain half-width	Volatility			Skewness			Kurtosis		
	10%	50%	80%	10%	50%	80%	10%	50%	80%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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 11: 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).

Noise	Volatility			Skewness			Kurtosis		
	1%	5%	10%	1%	5%	10%	1%	5%	10%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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 12: 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).

Spacing	Volatility			Skewness			Kurtosis		
	1%	2%	5%	1%	2%	5%	1%	2%	5%
Panel A: Standard Scenario									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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									
<i>Central Moments</i>									
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
<i>Quantile Moments</i>									
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