

Tails, Fears and Risk Premia*

Tim Bollerslev[†] and Viktor Todorov[‡]

First draft: January, 2009
This version: January 11, 2011

Abstract

We show that the compensation for rare events accounts for a large fraction of the average equity and variance risk premia. Exploiting the special structure of the jump tails and the pricing thereof we identify and estimate a new Investor Fears index. The index reveals large time-varying compensations for fears of disasters. Our empirical investigations involve new extreme value theory approximations and high-frequency intraday data for estimating the expected jump tails under the statistical probability measure, and short maturity out-of-the money options and new model-free implied variation measures for estimating the corresponding risk neutral expectations.

Keywords: rare events, jumps, high-frequency data, options, fears, extreme value theory, equity risk premium, variance risk premium.

JEL classification: C13, C14, G10, G12.

*We would like to thank Yacine Aït-Sahalia, Torben G. Andersen, Snehal Banerjee, Robert Barro, Geert Bekaert, Jules van Binsbergen, Oleg Bondarenko, Mikhail Chernov, Rama Cont, Dobrislav Dobrev, Darrell Duffie, Eric Engstrom, Jeremy Graveline, Friedrich Hubalek, Jean Jacod, Claudia Kluppelberg, Kai Li, Sydney Ludvigson, Lasse Pedersen, Mark Podolskij, Petra Posadel, Walter Schachermayer, George Tauchen and Amir Yaron, our discussants Alexandra Dias, Xavier Gabaix and Yuhang Xing, two anonymous referees, as well as seminar participants at EDHEC, the Federal Reserve Board, UIC, the June 2009 CREATES-Stevanovic Center conference in Skagen, the October 2009 SoFiE conference on “Liquidity, Credit Risk and Extreme Events” in Chicago, the January 2010 conference on the “Interplay between Financial and Insurance Mathematics, Statistics and Econometrics” at Wolfgang Pauli Institute, Vienna University, the 2010 Warwick Business School Conference on Derivatives, Volatility and Correlation, the June 2010 FERM symposium in Taipei, the 2010 WFA meetings and the 2010 NBER Fall Asset Pricing meeting for their comments and helpful discussions related to the paper. We would also like to thank OptionMetrics for providing us with the options data, Lai Xu for help with extracting and organizing the data, and Nicola Fusari for the help with the option pricing codes in the Monte Carlo. The research was partly funded by NSF grant SES-0957330 to the NBER. Bollerslev would also like to acknowledge the support of CREATES funded by the Danish National Research Foundation.

[†]Department of Economics, Duke University, Durham, NC 27708, and NBER and CREATES; e-mail: bollerv@duke.edu.

[‡]Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: v-todorov@kellogg.northwestern.edu

“So what are policymakers to do? First and foremost, reduce uncertainty. Do so by removing tail risks, and the perception of tail risks.”

Olivier Blanchard, chief economist, IMF, *The Economist*, January 31, 2009

The lack of investor confidence is frequently singled out as one of the main culprits behind the massive losses in market values in the advent of the Fall 2008 financial crises. At the time, the portrayal in the financial news media of different doomsday scenarios in which the stock market would have declined even further were also quite commonplace. Motivated by these observations, we provide a new theoretical framework for better understanding the way in which the market prices and perceives jump tail risks. Our estimates rely on the use of actual high-frequency intraday data and short maturity out-of-the-money options. Our empirical results based on data for the S&P 500 index spanning the period from 1990 until mid-2007 show that the market generally incorporates the possible occurrence of rare disasters in the way in which it prices risky payoffs, and that the fear of such events accounts for a surprisingly large fraction of the historically observed equity and variance risk premia. Extending our results through the end of 2008, we document an even larger role of investor fears during the recent financial market crises. As such, our findings implicitly support the above conjecture in that removing the perception of jump tail risk was indeed crucial in restoring asset values at the time.

More precisely, by exploiting the special structure of the jump tails and the pricing thereof we identify and estimate a new Investors Fears index.¹ Our identification of investors fears rests on the distinctly different economic roles played by the compensation for the pathwise variation in asset prices versus the compensation for the possible occurrence of “large” rare jump events. These two separate sources of variation have traditionally have been treated as latent constituents of the market risk premia. Meanwhile, our findings suggest that compensation for the former, which is naturally associated with temporal variation in the investment opportunity set, is rather modest, while the compensation attributable to the latter, and the fears of rare events, is both time-varying and often quite large.²

Our empirical estimates rely on a series of new statistical procedures for backing out the objective expectations of jump tail events and the market’s pricing thereof. In particular, building on new extreme value theory developed in Bollerslev and Todorov (2010), we use the relatively frequent “medium” size jumps in high-frequency intraday returns for reliably estimating the expected interdaily tail events under the statistical probability measure, thereby explicitly avoiding peso type problems in inferring the actual jump tails.

¹Our Investor Fears index is conceptually different from standard measures of tail risks, which only depend on the actual probabilities for the occurrences of tail events but not the pricing thereof; see, e.g., the discussion in Artzner et al. (1999). Most traditional tail risk measures, of course, also apply with any portfolio, while the Investor Fears index pertains explicitly to the pricing of jump tail risk in the aggregate market portfolio.

²The relatively minor role played by temporal variation in the investment opportunity set for explaining the risk premia suggested by our results is also consistent with the evidence in Chacko and Viceira (2005), which indicate only modest gains from standard intertemporal hedging in the context of dynamic consumption and portfolio choice.

Similarly, our estimates for the market risk-neutral expectations are based on actual short maturity out-of-the money options and the model-free variation measures originally utilized by Carr and Wu (2003), further developed here to recover the complete risk-neutral tail jump density. Intuitively, while the continuous price variation and the possibility of jumps both affect short maturity options, their relative importance varies across different strikes, which in turn allow us to separate the valuation of the two risks in a model-free manner.

The general idea behind the paper is related to an earlier literature that seek to explain the observed differences in the time series behavior of options prices and the prices of the underlying asset through the pricing of jump risk; e.g., Andersen et al. (2002), Bates (1996, 2000), Broadie et al. (2009a), Eraker (2004) and Pan (2002). The paper is also related to the work of Aït-Sahalia et al. (2001), who compare non-parametric estimates of the risk-neutral state price densities from time series of returns to the densities estimated directly from options prices. All of these earlier studies, however, rely on specific, typically affine, parametric stochastic volatility jump diffusion models, or an assumption of no unspanned stochastic volatility. Our approach is distinctly different in relying on flexible non-parametric procedures that is able to accommodate rather complex dynamic tail dependencies as well as time-varying stochastic volatility. Our focus on the aggregate market risk premia and the notion of investor fears implicit in the wedge between the estimated objective and risk-neutral jump tails, along with the richer data sources utilized in the estimation, also set the paper apart from this earlier literature.

Our empirical results suggest that on average close to five percent of the equity premium (in absolute terms) may be attributed to the compensation for rare disaster events. Related option-based estimates for the part of the equity premium due to compensation for jump risk have previously been reported in the literature by Broadie et al. (2009a), Eraker (2004), Pan (2002), and Santa-Clara and Yan (2010) among others. However, as noted above, these earlier studies are based on fairly tightly parameterized jump-diffusion models and restrictive assumptions about the shapes and dependencies in the tails.³ Counter to these assumptions, our new non-parametric approach reveals a series of intriguing dynamic dependencies in the tail probabilities, including time-varying jump intensities and a tendency for fatter tails in periods of high overall volatility. At the same time, our results also suggest that the risk premium for tail events cannot solely be explained by the level of the volatility.

In parallel to the equity premium, defined by the difference between the statistical and risk-neutral expectation of the aggregate market returns, the variance risk premium is naturally defined as the difference between the statistical and risk-neutral expectation of the corresponding forward variation.⁴ Taking our analysis one step further, we show that on average more than half of the historically observed variance risk premium is directly attributable to disaster risk, or differences in the jump tails of the risk-neutral and objective

³Anderson et al. (2003) and Broadie et al. (2009a) both forcefully argue from very different perspectives that misspecified parametric models can result in highly misleading estimates for the risk premia.

⁴Several recent studies have sought to explain the existence of a variance risk premium, and how the premium correlates with the underlying returns, within the context of various equilibrium based pricing models; see, e.g., Bakshi and Kapadia (2003), Bakshi and Madan (2006), Bollerslev et al. (2009), Carr and Wu (2009), Drechsler and Yaron (2011), Eraker (2008), and Todorov (2010).

distributions. Moreover, while the left and right jump tails in the statistical distribution are close to symmetric, the contribution to the overall risk-neutral variation coming from the left tail associated with dramatic market declines is several times larger than the variation attributable to the right tail. Therefore, even though the VIX volatility index formally contains compensation for different risks, i.e., time-varying volatilities and jump intensities as well as fears for jump tail events, our separation shows that a non-trivial portion of the index may be attributed to the latter component and notions of investor fears.⁵

At a more general level, our results hold the promise of a better understanding of the economics behind the historically large average equity and variance risk premia observed in the data. Most previous equilibrium based explanations put forth in the literature have effectively treated the two risk premia in isolation, relying on different mechanisms for mapping the dynamics of the underlying macroeconomic fundamentals into asset prices. Instead, our analysis shows that high-frequency aggregate market and derivative prices in combination contain sufficient information to jointly identify both premia in an essentially model-free manner, and that the fear of rear events play a very important role in determining both. As such, any satisfactory equilibrium based asset pricing model must be able to generate large and time-varying compensations for the possible occurrence of rare disasters.⁶

The plan for the rest of the paper is as follows. We begin in the next section with a discussion of the basic setup and assumptions, along with the relevant formal definitions of the equity and variance risk premia. Section 2 outlines the procedures that we use for separating jumps and continuous price variation under the risk-neutral distribution. Section 3 details the methods that we use for the comparable separation under the statistical measure. Section 4 discusses the data and our initial empirical results. Our main empirical findings related to the tail risks and the risk premia are given in Section 5. Section 6 reports on various robustness checks related to the dependencies in the tail measures and the accuracy of the asymptotic approximations underlying our results. Section 7 concludes. Most of the technical proofs are deferred to the Appendix, as are some of the details concerning our handling of the options and intraday data, as well as the econometric modeling procedures.

⁵As such, this does lend some credence to the common use of the term “investor’s fear gauge” as a moniker for the VIX volatility index, albeit a rather imperfect proxy at that; see, e.g., Whaley (2009). Instead, the new Investor Fears index developed here “cleanly” isolates the fear component.

⁶The idea that rare disasters, or tail events, may help explain the equity premium and other empirical puzzles in asset pricing finance dates back at least to Rietz (1988). While the original work by Rietz (1988) was based on a peso type explanation and probabilities of severe events that exceed those materialized in sample, Barro (2006) has more recently argued that when calibrated to international data, actual disaster events do appear frequent enough to meaningfully impact the size of the equity premium. Further building on these ideas, Gabaix (2010) and Wachter (2010) have also recently shown that explicitly incorporating time-varying risks of rare disasters in otherwise standard equilibrium based asset pricing models may help explain the apparent excess volatility of aggregate market returns.

1 Risk Premia and Jumps

The continuous-time no-arbitrage framework that underly our empirical investigations is very general and essentially model-free. It includes all parametric models previously analyzed and estimated in the literature as special cases. We begin with a discussion of the basic setup and notation.

1.1 Setup and Assumptions

To set up the notation, let F_t denote the futures price for the aggregate market portfolio expiring at some undetermined future date. Also, let the corresponding logarithmic price be denoted by $f_t \equiv \log(F_t)$. The absence of arbitrage, implies that the futures price should be a semimartingale.⁷ Restricting this just slightly by assuming an Itô semimartingale along with jumps of finite variation, the dynamics of the futures price may be expressed as,⁸

$$\frac{dF_t}{F_t} = \alpha_t dt + \sigma_t dW_t + \int_{\mathbb{R}} (e^x - 1) \tilde{\mu}(dt, dx), \quad (1.1)$$

where α_t and σ_t denote locally bounded, but otherwise unspecified drift and instantaneous volatility processes, respectively, and W_t is a standard Brownian motion. The last term accounts for any jumps, or discontinuities, in the price process through the so-called compensated jump measure $\tilde{\mu}(dt, dx) \equiv \mu(dt, dx) - \nu_t^{\mathbb{P}}(dx)dt$, where $\mu(dt, dx)$ is a simple counting measure for the jumps and $\nu_t^{\mathbb{P}}(dx)dt$ denotes the compensator (stochastic intensity) of the jumps for $\nu_t^{\mathbb{P}}(dx)$ being predictable and such that $\int_{\mathbb{R}} (x^2 \wedge 1) \nu_t^{\mathbb{P}}(dx)$ is locally integrable.⁹ This is a very general specification, and aside from the implicit weak integrability conditions required for the existence of some of the expressions discussed below, we do not make any further assumptions about the properties of the jumps, or the form of the stochastic volatility process.¹⁰

In order to better understand the impact and pricing of the two separate components, it is instructive to consider the total variation associated with the market price. Specifically, let $QV_{[t,T]}$ denote the quadratic variation of the log-price process over the $[t, T]$ time-interval,

$$QV_{[t,T]} \equiv \int_t^T \sigma_s^2 ds + \int_t^T \int_{\mathbb{R}} x^2 \mu(ds, dx). \quad (1.2)$$

The first term on the right-hand-side represents the variation attributable to the continuous-time stochastic volatility process σ_s ; i.e., the variation due to “small” price moves. The

⁷For the connection between the semimartingale assumption for the price and the notion of no arbitrage, see e.g., Delbaen and Schachermayer (1994).

⁸The assumption of finite variation for the jumps is inconsequential and for expositional purposes only. The “big” jumps are always of finite activity; i.e., only a finite number over any finite time interval.

⁹Intuitively, this renders the “demeaned” sum of the jumps $\int_{\mathbb{R}} (e^x - 1) \tilde{\mu}(dt, dx)$ a (local) martingale.

¹⁰The specification in (1.1) does exclude semimartingale processes whose characteristics are not absolutely continuous with respect to Lebesgue measure, e.g., certain time-changed Brownian motions where the time-change is a discontinuous process.

second term measures the variation due to jumps; i.e., “large” price moves. These risks are fundamentally different and present different challenges from a risk management perspective. While diffusive risks can be hedged/managed by continuously rebalancing a portfolio exposed to those risks, the locally unpredictable nature of jumps means that such a strategy will not work. Consequently, these two types of risks may also demand different compensation, as manifest in the form of different contributions to the aggregate equity and variance risk premia.

1.2 Equity and Variance Risk Premia

The equity risk premium represents the compensation directly associated with the uncertainty about the future price level. The variance risk premium refers to the compensation for the risk associated with temporal changes in the variation of the price level.¹¹

In particular, let \mathbb{Q} denote the risk-neutral distribution associated with the general dynamics in equation (1.1). The equity and variance risk premia are then formally defined by,¹²

$$ERP_t \equiv \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \left(\frac{F_T - F_t}{F_t} \right) - \mathbb{E}_t^{\mathbb{Q}} \left(\frac{F_T - F_t}{F_t} \right) \right), \quad (1.3)$$

and

$$VRP_t \equiv \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} (QV_{[t,T]}) - \mathbb{E}_t^{\mathbb{Q}} (QV_{[t,T]}) \right), \quad (1.4)$$

corresponding to the expected excess return on the market and the expected payoff on a (long) variance swap on the market portfolio, respectively. Empirically, of course, the sample equity risk premium is generally positive and historically “large”, while the sample variance risk premium as defined above is on average negative and equally “large” and puzzling.¹³

¹¹Investors desire to hedge against intertemporal shifts in riskiness has led to the recent advent of many new financial instruments with their payoffs directly tied to various notions of realized price variation. Especially prominent among these are variance swap, or forward, contracts on future realized variances. Exotic so-called gap options explicitly designed to hedge against large price moves over short daily time intervals have also recently been introduced on the OTC market.

¹²From here on we will suppress the dependence on the horizon T for notational convenience. Also, the discrete-time equity premium defined here should not be confused with the instantaneous premium, or drift, defined in equation (1.5) below. The latter serves a particular convenient role in explicitly separating the different types of risks, as manifest in equation (1.12).

¹³The equity premium puzzle and the failure of standard consumption-based asset pricing models to explain the sample ERP_t is arguably one of the most studied issues in academic finance over the past two decades; see, e.g., the discussion in Campbell and Cochrane (1999) and Bansal and Yaron (2004) and the many references therein. Representative agent equilibrium models involving time-separable utility also rule out priced volatility risk and corresponding hedging demands, and in turn imply that $VRP_t = 0$. Recent rational equilibrium based pricing models designed to help explain the positive variance risk premium by incorporating jumps and/or time varying volatility-of-volatility in the consumption growth process of the representative agent include Bollerslev et al. (2009), Drechsler and Yaron (2011) and Eraker (2008).

In order to disentangle the distributional features that account for these well-documented differences in the \mathbb{P} and \mathbb{Q} expectations, let $\nu_t^{\mathbb{Q}}(dx)dt$ denote the compensator, or intensity, for the jumps under \mathbb{Q} (recall our absolute continuity assumption for the \mathbb{P} jump compensator and the fact that the probability measures \mathbb{P} and \mathbb{Q} are equivalent). Also, let λ_t denote the drift that turns W_t into a Brownian motion under \mathbb{Q} ; i.e., $W_t^{\mathbb{Q}} = W_t + \int_0^t \lambda_s ds$.¹⁴ The drift term for the futures price in the \mathbb{P} distribution in equation (1.1) must therefore satisfy,

$$\alpha_t = \lambda_t \sigma_t + \int_{\mathbb{R}} (e^x - 1) \nu_t^{\mathbb{P}}(dx) - \int_{\mathbb{R}} (e^x - 1) \nu_t^{\mathbb{Q}}(dx). \quad (1.5)$$

This general expression directly manifests the compensation that is required by investors to accept the different types of equity price related risks.

Utilizing this result, the equity risk premium is naturally decomposed as,

$$ERP_t = ERP_t^c + ERP_t^d, \quad (1.6)$$

where

$$ERP_t^c \equiv \frac{1}{T-t} \mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \frac{F_s}{F_t} \lambda_s \sigma_s ds \right), \quad (1.7)$$

and

$$ERP_t^d \equiv \frac{1}{T-t} \mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \int_{\mathbb{R}} \frac{F_s}{F_t} (e^x - 1) \nu_s^{\mathbb{P}}(dx) ds - \int_t^T \int_{\mathbb{R}} \frac{F_s}{F_t} (e^x - 1) \nu_s^{\mathbb{Q}}(dx) ds \right), \quad (1.8)$$

and each of the two terms represent the unique contribution to the premium coming from diffusive and jump risk, respectively.

Similarly, the variance risk premium may be expressed as,

$$VRP_t = VRP_t^c + VRP_t^d, \quad (1.9)$$

where

$$VRP_t^c = \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \sigma_s^2 ds \right) - \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \sigma_s^2 ds \right) \right), \quad (1.10)$$

and

$$VRP_t^d = \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \int_{\mathbb{R}} x^2 \nu_s^{\mathbb{P}}(dx) ds \right) - \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{\mathbb{R}} x^2 \nu_s^{\mathbb{Q}}(dx) ds \right) \right), \quad (1.11)$$

represent the diffusive and jump risk components, respectively. As equations (1.10)-(1.11) make clear, the presence of jumps may result in a non-zero premium even with i.i.d. returns.¹⁵

¹⁴Standard asset pricing theory implies that all futures (and forward) contracts must be martingales under the risk-neutral \mathbb{Q} measure; see, e.g., Duffie (2001). This result is subject to a boundedness condition on the interest rate process, but such assumption can be further relaxed; see Pozdnyakov and Steele (2004).

¹⁵It is possible to further separate VRP_t^c into two components - diffusive and jump-type changes in the stochastic volatility. We purposely do not explore this, however, as our main focus remains on the separation of the fears of disasters from the premia attached to changes in the investment opportunity set, and the type of changes underlying the evolution in σ_t has no effect on that separation. Further decomposing VRP_t^c would also necessitate stronger parametric assumptions than the ones employed here.

The jump risk links directly the equity and variance risk premium. By contrast, the price of diffusive price risk does not enter the variance risk premium. Previous attempts in the literature at trying to separately identify and estimate the risk premia due to jumps have been based on explicit and somewhat restrictive parametric models, in turn resulting in different and sometimes conflicting numerical findings; see, e.g., the discussion in chapter 15 of Singleton (2006). Meanwhile, as equations (1.6) and (1.9) make clear, the compensation for jumps will generally affect the equity and variance risk premia in distinctly different, and statistically readily identifiable, ways.

In particular, using actual high-frequency data and short maturity options it is possible to non-parametrically estimate,¹⁶

$$ERP_t(k) = \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \int_{|x|>k} (e^x - 1) \nu_s^{\mathbb{P}}(dx) ds \right) - \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{|x|>k} (e^x - 1) \nu_s^{\mathbb{Q}}(dx) ds \right) \right), \quad (1.12)$$

and

$$VRP_t(k) = \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \int_{|x|>k} x^2 \nu_s^{\mathbb{P}}(dx) ds \right) - \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{|x|>k} x^2 \nu_s^{\mathbb{Q}}(dx) ds \right) \right), \quad (1.13)$$

where $k > 0$ refers to the threshold that separates “large” from “small” jumps.¹⁷ Comparing these estimates with the corresponding estimates for ERP_t with VRP_t , we may therefore gauge the importance of fears and rare events in explaining the equity and variance risk premia from directly observable financial data quite generally. We begin our discussion of the relevant empirical procedures with the way in which we quantify the risk-neutral \mathbb{Q} measures.

2 Risk-Neutral Measures

There is long history of using options prices to infer market expectations and corresponding options implied volatilities. Traditionally, these estimates have been based on specific assumptions about the underlying \mathbb{P} distribution and the pricing of systematic risk(s). However, as emphasized in the literature more recently, the risk-neutral expectation of the, normalized by the horizon $T - t$, total quadratic variation of the price process in equation (1.2),

$$QV_t^{\mathbb{Q}} \equiv \frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \sigma_s^2 ds \right) + \frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{\mathbb{R}} x^2 \mu(ds, dx) \right), \quad (2.1)$$

¹⁶The integrands in equations (1.8) and (1.12) obviously differ by F_s/F_t . Alternatively, we could have defined the equity risk premium in terms of logarithmic returns, in which case F_s/F_t would not appear in equation (1.8). In that case, however, the basic definition would instead involve convexity adjustment terms.

¹⁷The choice of k and other practical implementation issues are discussed further in the Appendix below.

may be estimated in a completely model-free manner by a portfolio of options with different strikes; see, e.g., Bakshi et al. (2003), Britten-Jones and Neuberger (2000) and Carr and Wu (2009).¹⁸ Meanwhile, as the expression in equation (2.1) makes clear, and in parallel to the discussion above, the expected quadratic variation under \mathbb{Q} is naturally decomposed into the expected variation associated with diffusive, or “small” price moves, and discontinuous, or “large” price moves. Building on these insights, we show below how options prices may be used in estimating these two separate parts of the risk-neutral distribution in an essentially model-free fashion.¹⁹

2.1 Jumps and Tails

Our estimates for the jump risk perceived by investors are based on close-to-maturity deep out-of-the-money options. The idea of using deep out-of-the-money options to more effectively isolate jump risk has previously been used in the estimation of specific parametric affine jump-diffusion stochastic volatility models by Bakshi et al. (1997), Bates (1996, 2000), Pan (2002) and Eraker (2004) among others, while Pan and Liu (2003) and Pan et al. (2003) have considered similar fully parametric approaches in the context of hedging portfolio jump risks.²⁰ In contrast to these studies, our estimates for the jump risk premia are entirely non-parametric.

In particular, building on the ideas in Carr and Wu (2003), who used the different rate of decay of the diffusive and jump risk for options with various levels of moneyness as a way to test for jumps, we rely on short-maturity options to actually identify the compensation for jump risk. Intuitively, short-maturity out-of-the-money options will remain worthless unless a rare event, or a “big” jump, occurs before expiration. As such, this allows us to infer the relevant risk-neutral expectations through appropriately scaled option prices.

Specifically, let $C_t(K)$ and $P_t(K)$ denote the price of a call option, respectively put option, with a strike of K recorded at time t expiring at some future date, T . It follows that for

¹⁸This also mirrors the way in which the CBOE now calculates the VIX volatility index for the S&P 500. The approximation used by the CBOE formally differs from the expression in equation (2.1) by $\frac{1}{T-t}\mathbb{E}_t^{\mathbb{Q}}\left(\int_t^T\int_{\mathbb{R}}2\left(e^x-1-x-\frac{x^2}{2}\right)\mu(ds,dx)\right)$; for further discussion of the approximation errors involved in this calculation see the *white paper* on the CBOE website, as well as Carr and Wu (2009) and Jiang and Tian (2005, 2007).

¹⁹Previous attempts at non-parametrically estimating the risk-neutral distribution implied by options prices include Aït-Sahalia and Lo (1998, 2000), Bakshi et al. (2003) and Rosenberg and Engle (2002) among others. Aït-Sahalia et al. (2001), in particular, have previously sought to quantify the amount of additional jump risk needed to make the \mathbb{P} tails comparable to the options implied \mathbb{Q} tails, by calibrating the jump frequency to match the estimated skewness (or kurtosis) of the state price density. To the best of our knowledge, however, none of these studies have sought to estimate the portion of the two risk premia that are directly attributable to the jump tails and the fears thereof.

²⁰As previously noted, the recent exotic so-called gap options, or gap risk swaps, are explicitly designed to hedge against gap events, or jumps, that is a one-day large price move in the underlying; see Cont and Tankov (2009) and Tankov (2009) for further discussion on the pricing of these contracts using different parametric models as well as their hedging using standard exchanged-traded out-of-the-money options.

$T \downarrow t$ and $K > F_{t-} \equiv \lim_{s \uparrow t} F_s$, where \lim denotes the limit from the left,

$$e^{r(t,T)}C_t(K) \approx \int_t^T \mathbb{E}_t^{\mathbb{Q}} \left(\int_{\mathbb{R}} 1_{\{F_{s-} < K\}} (F_{s-}e^x - K)^+ \nu_s^{\mathbb{Q}}(dx) \right) ds, \quad (2.2)$$

while for $K < F_{t-}$,

$$e^{r(t,T)}P_t(K) \approx \int_t^T \mathbb{E}_t^{\mathbb{Q}} \left(\int_{\mathbb{R}} 1_{\{F_{s-} > K\}} (K - F_{s-}e^x)^+ \nu_s^{\mathbb{Q}}(dx) \right) ds, \quad (2.3)$$

where $x^+ \equiv \max\{0, x\}$. Importantly, the expressions on the right-hand sides of (2.2) and (2.3) involve only the jump measure. Over short time intervals changes in the price due to the continuous component are invariably small relative to the possible impact of “large” jumps, and the diffusive part may simply be ignored. Moreover, for $T \downarrow t$ there will be at most one “large” jump. The right-hand sides of equations (2.2) and (2.3) represent the conditional expected values of the two different option payoffs under these approximations. The payoffs will be zero if no “large” jump occurred and otherwise equal to the difference between the price after the jump and the strike price of the option. The proof of Proposition 1 below given in Appendix A contains a more formal analysis of the errors involved in the approximations and their orders of magnitude.

Going one step further, the expressions in equations (2.2) and (2.3) may be used in the construction of the following model-free risk-neutral jump tail measures,

$$RT_t^{\mathbb{Q}}(k) \equiv \frac{1}{T-t} \int_t^T \int_{\mathbb{R}} (e^x - e^k)^+ \mathbb{E}_t^{\mathbb{Q}}(\nu_s^{\mathbb{Q}}(dx)) ds \approx \frac{e^{r(t,T)}C_t(K)}{(T-t)F_{t-}}, \quad (2.4)$$

and

$$LT_t^{\mathbb{Q}}(k) \equiv \frac{1}{T-t} \int_t^T \int_{\mathbb{R}} (e^k - e^x)^+ \mathbb{E}_t^{\mathbb{Q}}(\nu_s^{\mathbb{Q}}(dx)) ds \approx \frac{e^{r(t,T)}P_t(K)}{(T-t)F_{t-}}, \quad (2.5)$$

where $k \equiv \ln(\frac{K}{F_{t-}})$ denotes the log-moneyness. Unlike the variance replicating portfolio used in the construction of the VIX index and the approximation to $QV_t^{\mathbb{Q}}$ in equation (2.1), which holds true for any value of T , the approximations for $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ in equations (2.4) and (2.5) depend on $T \downarrow t$.

In practice, of course, we invariably work with close-to-maturity options and a finite horizon $T - t$, so the effect of the diffusive volatility might be non-trivial. In Section 6, we report representative results from an extensive Monte Carlo simulation study designed to investigate the “finite sample” accuracy of the approximations across different models, option strikes K , and time-to-maturity $T - t$. Taken as whole, these results suggest that the error involved in approximating $LT_t^{\mathbb{Q}}$ through equation (2.5) is in fact quite trivial for options and empirical settings designed to mimic that of the actual data. The estimation error for $RT_t^{\mathbb{Q}}$ in equation (2.4) tend to be somewhat larger, albeit still fairly small in an absolute sense. Intuitively, investors primarily fear “large” negative price moves, making the

negative jumps in the risk-neutral distribution more pronounced and easier to separate from the diffusive component.

Although the $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ measures are related to the tail jump density, they do not directly correspond to the risk-neutral expectations required for the calculation of the $ERP_t(k)$ and $VRP_t(k)$ jump premia. In order to get at these and the risk-neutral tail density, we rely on an additional Extreme Value Theory (EVT) type approximation.²¹ This approximation allows us to easily extrapolate the tail behavior from the estimates of $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ for a range of different values of k . The only additional assumption required for these calculations concerns the format of the risk-neutral jump density.

Specifically, we will assume that

$$\nu_t^{\mathbb{Q}}(dx) = (\varphi_t^+ 1_{\{x>0\}} + \varphi_t^- 1_{\{x<0\}}) \nu^{\mathbb{Q}}(x) dx, \quad (2.6)$$

where the unspecified φ_t^{\pm} stochastic processes allow for temporal variation in the jump arrivals and $\nu^{\mathbb{Q}}(x)$ denotes any valid Lévy density. This is a very weak assumption on the general form of the tails, merely restricting the time-variation to be the same across different jump sizes. This assumption is trivially satisfied by all of the previously estimated parametric jump-diffusion models referred to above, most of which restrict $\varphi_t^{\pm} \equiv 1$, or $\varphi_t^{\pm} \equiv \sigma_t^2$.

The following Proposition, the proof of which is deferred to Appendix A, formalizes the results underlying our estimation of $\mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{|x|>k} (e^x - 1) \nu_s^{\mathbb{Q}}(dx) ds \right)$ and $\mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{|x|>k} x^2 \nu_s^{\mathbb{Q}}(dx) ds \right)$ from $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$. Intuitively, by tracking the slope of the latter for increasingly deeper out-of-the-money options, it is possible to infer the tail behavior of $\nu_t^{\mathbb{Q}}(dx)$.²²

Proposition 1 *Assume that F_t and $\nu_t^{\mathbb{Q}}(dx)$ satisfy equations (1.1) and (2.6), respectively. Denote $\psi^+(x) = e^x - 1$ and $\nu_{\psi}^{\mathbb{Q}+}(x) = \frac{\nu^{\mathbb{Q}}(\ln(x+1))}{x+1}$ for $x > 0$, and $\psi^-(x) = e^{-x}$ and $\nu_{\psi}^{\mathbb{Q}-}(x) = \frac{\nu^{\mathbb{Q}}(-\ln(x))}{x}$ for $x < 0$. Define $\bar{\nu}_{\psi}^{\mathbb{Q}\pm}(x) = \int_x^{\infty} \nu_{\psi}^{\mathbb{Q}\pm}(u) du$, and assume that $\bar{\nu}_{\psi}^{\mathbb{Q}\pm}(x) = x^{-p_{\mathbb{Q}}^{\pm}} L^{\pm}(x)$ for $p_{\mathbb{Q}}^{\pm} > 0$, where $L^{\pm}(x)$ satisfy $L^{\pm}(tx)/L^{\pm}(x) = 1 + O(\tau^{\pm}(x))$ as $x \uparrow \infty$ for each $t > 0$, where $\tau^{\pm}(x) > 0$, $\tau^{\pm}(x) \rightarrow 0$ as $x \uparrow \infty$, and $\tau^{\pm}(x)$ are nonincreasing. Then for arbitrary $t > 0$, $T - t \downarrow 0$, $K_1 \uparrow \infty$, and $K_2 \uparrow \infty$, with $K_2/K_1 > 1$ constant,*

$$\log \left(\frac{RT_t^{\mathbb{Q}}(K_1/F_{t-})}{RT_t^{\mathbb{Q}}(K_2/F_{t-})} \right) \xrightarrow{\mathbb{P}} (p_{\mathbb{Q}}^+ - 1) \log \left(\frac{K_2 - F_{t-}}{K_1 - F_{t-}} \right), \quad (2.7)$$

while for $K_1 \downarrow 0$ and $K_2 \downarrow 0$ with $K_1/K_2 > 1$ constant,

$$\log \left(\frac{LT_t^{\mathbb{Q}}(K_1/F_{t-})}{LT_t^{\mathbb{Q}}(K_2/F_{t-})} \right) \xrightarrow{\mathbb{P}} (p_{\mathbb{Q}}^- + 1) \log \left(\frac{K_1}{K_2} \right). \quad (2.8)$$

The risk premia and the $ERP_t(k)$ and $VRP_t(k)$ measures, of course, also depend on the jump tails in the statistical distribution. We next discuss the high-frequency realized variation measures and reduced form forecasting models that we rely on in estimating these.

²¹For a general discussion of Extreme Value Theory see, e.g., Embrechts et al. (2001).

²²Our actual estimation is slightly more general than the result of Proposition 1 in the sense that it doesn't restrict a scaling parameter; the details are in Appendix A.

3 Objective Measures

Our estimates of the relevant expectations under the objective, or statistical, probability measure \mathbb{P} are based on high-frequency intraday data and corresponding realized variation measures for the diffusive and dis-continuous components of the total variation in equation (1.2). In addition, we rely on newly developed EVT-based approximations and non-parametric reduced form modeling procedures for translating the ex-post measurements into forward looking \mathbb{P} counterparts to the $QV_t^{\mathbb{Q}}$, $RT_t^{\mathbb{Q}}$, and $LT_t^{\mathbb{Q}}$ expectations defined above. The foundations for the EVT approximations are formally based on the general theoretical results in Bollerslev and Todorov (2010), and we provide details on the calculations involved in these approximations in the Appendix. We continue the discussion here with a description of our high-frequency based realized variation measures.

3.1 Realized Jumps and Volatility

For ease of notation, we normalize the unit time-interval to represent a day, where by convention the day starts with the close of trading on the previous day. The resulting daily $[t, t + 1]$ time interval is naturally broken into the $[t, t + \pi_t]$ overnight time period and the $[t + \pi_t, t + 1]$ active part of the trading day. The opening time of the market $t + \pi_t$ is obviously fixed in calendar time. Conceptually, however, it is useful to treat $\{\pi_t\}_{t=1,2,\dots}$ as a latent discrete-time stochastic process that implicitly defines the proportion of the total day t variation due to the overnight price change. With n equally spaced high-frequency price observations, this leaves us with a total of $n - 1$ intraday increments for day $t + 1$, each spanning a time interval of length $\Delta_{n,t} \equiv \frac{1-\pi_t}{n}$, say $\Delta_i^{n,t} f \equiv f_{t+\pi_t+i\Delta_{n,t}} - f_{t+\pi_t+(i-1)\Delta_{n,t}}$ for $i = 1, \dots, n - 1$.²³

The realized variation for day t , denoted RV_t , is simply defined as the summation of the squared high-frequency increments over the active part of the trading day. The realized variation concept was first introduced into the economics and finance literatures by Andersen and Bollerslev (1998), Andersen et al. (2002), and Barndorff-Nielsen and Shephard (2001). The key insight stems from the basic result that for increasingly finer sampled increments, or $n \uparrow \infty$, RV_t consistently estimates the total ex-post variation of the price process. Formally,

$$RV_t \equiv \sum_{i=1}^{n-1} (\Delta_i^{n,t} f)^2 \xrightarrow{\mathbb{P}} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds + \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx), \quad (3.1)$$

where the right-hand-side equals the quadratic variation of the price over the $[t + \pi_t, t + 1]$ time interval, as defined by the corresponding increment to the total variation in equation (1.2).

As discussed above, the quadratic variation obviously consists of two distinct components: the variation due to the continuous sample price path and the variation coming from jumps.

²³Even though it is convenient to treat the sampling times $(t + \pi_t, t + \pi_t + \Delta_{n,t}, \dots, t + 1)$ as stochastic, they are obviously fixed in calendar time. As discussed further in the Appendix, all we need from a theoretical perspective is for the value of π_t to be (conditionally) independent of the f_t logarithmic price process.

It is possible to separately estimate these two components by decomposing the summation of the squared increments into separate summations of the “small” and “large” price changes, respectively. In particular, it follows that for the continuous variation

$$CV_t \equiv \sum_{i=1}^{n-1} (\Delta_i^{n,t} f)^2 1_{\{|\Delta_i^{n,t} f| \leq \alpha \Delta_{n,t}^{\varpi}\}} \xrightarrow{\mathbb{P}} \int_{t+\pi_t}^{t+1} \sigma_s^2 ds, \quad (3.2)$$

while for the jump variation

$$JV_t \equiv \sum_{i=1}^{n-1} (\Delta_i^{n,t} f)^2 1_{\{|\Delta_i^{n,t} f| > \alpha \Delta_{n,t}^{\varpi}\}} \xrightarrow{\mathbb{P}} \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 \mu(ds, dx), \quad (3.3)$$

where $\alpha > 0$ and $\varpi \in (0, 0.5)$ denote positive constants.²⁴ These measures for the continuous and jump variation were first implemented in the context of high-frequency financial data by Mancini (2001). They are based on the intuitive idea that over short time intervals, the continuous component of the price will be normally distributed with variance proportional to the length of the time interval. Thus, if the high-frequency price increment exceeds a certain threshold, it must be associated with a jump in the price.

Building on this same idea, the JV_t jump component may be further split into the variation coming from positive, or right tail, jumps,

$$RJV_t \equiv \sum_{i=1}^{n-1} |\Delta_i^{n,t} f|^2 1_{\{\Delta_i^{n,t} f > \alpha \Delta_{n,t}^{\varpi}\}} \xrightarrow{\mathbb{P}} \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 1_{\{x > 0\}} \mu(ds, dx), \quad (3.4)$$

and negative, or left tail, jumps

$$LJV_t \equiv \sum_{i=1}^{n-1} |\Delta_i^{n,t} f|^2 1_{\{\Delta_i^{n,t} f < -\alpha \Delta_{n,t}^{\varpi}\}} \xrightarrow{\mathbb{P}} \int_{t+\pi_t}^{t+1} \int_{\mathbb{R}} x^2 1_{\{x < 0\}} \mu(ds, dx). \quad (3.5)$$

It follows, of course, by definition that $RV_t = CV_t + RJV_t + LJV_t$.

The different ex-post variation measures discussed above all pertain to the actual realized price process. By contrast, the options implied diffusive and jump tail measures discussed in Section 2 provide ex-ante estimates of the risk-neutral expectations of the relevant sources of risks. We next discuss our translation of the actual realized risks into comparable measures of objective, or statistical, ex-ante expectations of these risks.

3.2 Expected Jumps and Tails

The most challenging problem in quantifying the statistical expectations in the equity and variance risk premia in equations (1.6) and (1.9), respectively, concerns the \mathbb{P} equivalent to the previously discussed $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ risk-neutral tail measures, say $RT_t^{\mathbb{P}}(k)$ and

²⁴The actual choice of α and ϖ , together with our normalization procedure designed to account for the strong intraday pattern in the volatility, are discussed in Appendix B.

$LT_t^{\mathbb{P}}(k)$, respectively. These measures pertain to the truly “large” jumps. However, these types of tail events are invariably rare, and possibly even non-existent, over a limited calendar time span. This is akin to a Peso type problem, and as such straightforward estimation of the tail expectations based on sample averages will generally not be reliable. Instead, we rely on EVT and the more frequent “medium” sized jumps implicit in the high-frequency intraday data to meaningfully extrapolate the behavior of the “large” jumps, and more precisely estimate the extreme parts of the jump tails.

Compared to conventional fully parametric procedures hitherto used in the literature for estimating the expected impact of jumps, our approach is much more flexible and has the distinct advantage of not relying on a tight connection between the jumps of all different sizes. Indeed, as discussed further below the statistical properties of the “small” and “large” jumps differ quite dramatically, as indirectly evidenced by the markedly different estimates for the tails obtained across competing parametric specifications that utilize *all* of the jumps for the tail estimation.

A completely general non-parametric approach allowing jumps of different sizes to exhibit arbitrary temporal variation is obviously impossible. Instead, we merely assume that the temporal variation in the jump intensity is a linear function of the stochastic volatility, as in²⁵

$$\nu_t^{\mathbb{P}}(dx) = (\alpha_0^- 1_{\{x < 0\}} + \alpha_0^+ 1_{\{x > 0\}} + (\alpha_1^- 1_{\{x < 0\}} + \alpha_1^+ 1_{\{x > 0\}}) \sigma_t^2) \nu^{\mathbb{P}}(x) dx, \quad (3.6)$$

where $\nu^{\mathbb{P}}(x)$ is a valid Lévy density, and α_0^{\pm} and α_1^{\pm} are nonnegative free parameters. This semiparametric structure does admittedly restrict the form of the temporal dependencies somewhat. Importantly, however, the assumption in equation (3.6) is only needed to pin down the dynamics of the risk premia. All of the unconditional results reported in the paper, i.e., the level of the risk premia, continue to hold true under the analogous weaker assumption for the \mathbb{Q} jump tails given in equation (2.6).²⁶ The assumption in equation (3.6) still provides a very flexible specification, and we provide direct support for its empirical validity in Section 6 below. Virtually all parametric models hitherto estimated in the literature have been based on simplified versions of this $\nu_t^{\mathbb{P}}(dx)$ measure, typically restricting $\alpha_1^{\pm} = 0$. By contrast, the more general specification in equation (3.6) explicitly allows for time-varying jump intensities as a function of the latent volatility process σ_t^2 . This includes so-called “self-affecting” jumps in which $\alpha_1^{\pm} > 0$, and σ_t^2 depends on past realized jumps, as suggested by the data. Of course, the specification also allows for the relative magnitudes of the constant and time-varying intensities to differ across positive and negative jumps.

Relying on (3.6) for describing the temporal dependencies in the jump intensity, the equivalents to the $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ measures under the statistical distribution, may be

²⁵This parallels the significantly weaker assumption for the risk-neutral jump intensity in equation (2.6). Extensions to allow for nonlinear functions of the volatility, or other directly observable state variables, are straightforward.

²⁶Also, given the differences in orders of magnitude of the \mathbb{P} and \mathbb{Q} tails reported in the empirical part, the practical import of assumption (3.6) for the actual empirical results is rather limited.

expressed as

$$RT_t^{\mathbb{P}}(k) = \left(\alpha_0^+ + \frac{\alpha_1^+}{T-t} \mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \sigma_s^2 ds \right) \right) \int_{\mathbb{R}} (e^x - e^k)^+ \nu^{\mathbb{P}}(x) dx, \quad (3.7)$$

and

$$LT_t^{\mathbb{P}}(k) = \left(\alpha_0^- + \frac{\alpha_1^-}{T-t} \mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \sigma_s^2 ds \right) \right) \int_{\mathbb{R}} (e^k - e^x)^+ \nu^{\mathbb{P}}(x) dx. \quad (3.8)$$

In order to estimate $RT_t^{\mathbb{P}}(k)$ and $LT_t^{\mathbb{P}}(k)$, and in turn $QV_t^{\mathbb{P}}$, we therefore only need to quantify the integrals with respect to the Lévy density $\nu^{\mathbb{P}}(x)$ and the conditional expectation of the integrated volatility $\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \sigma_s^2 ds \right)$.

In parallel to our estimation of the \mathbb{Q} jump tails, our estimates of the integrals with respect to $\nu^{\mathbb{P}}(x)$ are again based on the “medium” sized jumps coupled with EVT for meaningfully extrapolating the behavior of the tails. The separation of the jump intensity into the constant and time-varying part, i.e., α_0^{\pm} and α_1^{\pm} , respectively, is effectively achieved through the correlation between the occurrence of jumps and past volatility, as measured directly by the realized continuous variation CV_t in equation (3.2). A more detailed description of the relevant calculations are again given in the Appendix.

Our estimates for the expected integrated volatility, and the time-varying parts of $RT_t^{\mathbb{P}}$, $LT_t^{\mathbb{P}}$ and $QV_t^{\mathbb{P}}$, are based on reduced form time series modeling procedures, as originally advocated by Andersen et al. (2003). This provides a particular convenient and parsimonious approach for capturing the dynamic dependencies in the integrated volatility process and corresponding forecasts. From a statistical perspective, this is also a much easier problem than that of directly estimating the tails, as we have plenty of reliable data and previous empirical insights to go by.²⁷

In particular, it is well established that integrated volatility is a stationary, but highly persistent, process; see, e.g., Bollerslev et al. (2009) and the reduced form model estimates reported therein. It has also previously been documented that jumps in the price tend to be accompanied by jumps in the spot volatility σ_t , and that negative and positive jumps impact the volatility differently; i.e., a leverage effect through jumps. In order to accommodate all of these non-linear dependencies in the construction of our estimates for the expected integrated volatility, we formulate a Vector Autoregression (VAR) of order 22, or one month, for the four-dimensional vector consisting of CV_t , RJV_t , LJV_t , and the squared overnight increments $(f_t - f_{t+\pi_t})^2$. To keep the model tractable, we follow the HAR-RV modeling approach of Corsi (2009) and Andersen et al. (2007) and restrict all but the daily, weekly, and monthly lags, i.e., lags 1, 5, and 22, respectively, to be equal to zero. The forecasts for $\mathbb{E}_t^{\mathbb{P}} \left(\int_t^T \sigma_s^2 ds \right)$ is then simply obtained by recursively iterating the VAR to the relevant horizon. Further details concerning the model estimates and forecast construction are given in Appendix C.

²⁷The theoretical results in Andersen et al. (2004) and Sizova (2009) also indicate that the loss in forecast efficiency from reduced form procedures is typically small relative to fully efficient parametric procedures, and that the reduced form procedures have the added advantage of a build-in robustness to model misspecification.

Before turning to a discussion of the actual empirical results related to the \mathbb{P} and \mathbb{Q} tail measures and risk premia, we briefly discuss the data underlying our estimates.

4 Data and Preliminary Estimates

4.1 High-Frequency Data and Statistical Measures

Our high-frequency data for the S&P 500 futures contract was obtained from Tick Data Inc., and spans the period from January, 1990 to December, 2008. The prices are recorded at five-minute intervals, with the first price for the day at 8: 35 (CST) until the last price at 15: 15. All-in-all, this leaves us with a total of 81 intraday price observations for each of the 4,751 trading days in the sample.

The resulting realized variation measures all exhibit the well known volatility clustering effects, with a sharp increase in their levels toward the end of the sample and the Fall 2008 financial market crises. The realized variation measures in turn form the basis for our non-parametric estimates of the statistical expectations of the total quadratic variation and the left and right tail measures, $QV_t^{\mathbb{P}}$, $LT_t^{\mathbb{P}}$ and $RT_t^{\mathbb{P}}$. We will not discuss the details of the estimation results here, instead referring to the Appendix and the actual estimates reported therein. Nonetheless, it is instructive to highlight a few of the key empirical findings.

Directly in line with the discussion in the previous section, the stochastic volatility process is highly persistent. In addition, negative “realized” jumps significantly impact future volatility. In contrast to the strong own dynamic dependencies in the volatility, the effect of negative jumps is relatively short-lived, however. Our results also indicate that positive jumps do not affect future volatility. These findings are consistent with the empirical results in Barndorff-Nielsen et al. (2010) and Todorov and Tauchen (2011). They also support a two-factor stochastic volatility model, as in, e.g., Todorov (2010), in which one of the factors is highly persistent, while the other factor is short lived and mainly triggered by (negative) jumps.

Our estimation results also point to the existence of non-trivial temporal variation in the tail jump intensities, as manifest by the highly statistically significant estimates for α_1^{\pm} . Interestingly, the right and left jump tails behave quite similar under the statistical distribution with almost identical tail decays. Taken as a whole, however, our empirical results suggest much more complicated dependencies than those allowed by the parametric models hitherto estimated in the literature, further underscoring the advantages of our flexible non-parametric approach for characterizing the behavior of the \mathbb{P} jump tails.

4.2 Options Data and Risk-Neutral Measures

Our options data are obtained from OptionMetrics. The data consists of closing bid and ask quotes for all S&P 500 options traded on the CBOE. The data spans the period from January 1996 to the end of December 2008, for a total of 3,237 trading days. To avoid

obvious inconsistencies and problems with the price quotes, we apply similar filters to the ones used by Carr and Wu (2003, 2009) for “cleaning” the data.

The estimates for the two tail measures $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ are based on the closest to maturity options for day t with at least eight calendar days to expiration.²⁸ More specifically, we first calculate Black-Scholes implied volatilities for all out-of-the-money options with moneyness k using linear interpolation of options that bracket the targeted strike.²⁹ From these implied volatilities we then compute an out-of-the money option price with moneyness k , and in turn the $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ measures as described in Section 2 above.

Our calculation of the implied total variance measure $QV_t^{\mathbb{Q}}$, follows Carr and Wu (2009) and is similar to the approach used by the CBOE in their calculation of the VIX index. There are two differences, however, between our construction of $QV_t^{\mathbb{Q}}$ and the calculation of the VIX index, as described in the *white paper* available on the CBOE website. First, we rely on trading, or business, time, while the VIX index is constructed using a calendar time convention. Second, for compatibility with the tail measures $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$, we use the same closest to maturity options in the calculation of $QV_t^{\mathbb{Q}}$, whereas the VIX is based on linear interpolation of options that bracket a one-month maturity.

The three resulting risk-neutral variation measures are plotted in Figure 1. All of the measures are reported in annualized percentage form. Guided by the results from the Monte Carlo simulation study designed to investigate the accuracy of the approximations in equations (2.4) and (2.5) discussed in more detail in Section 6 below, here and throughout the rest of the analysis, the moneyness for the left and right tails are fixed at 0.9 and 1.1, respectively. As is immediately evident from Figure 1, the magnitude of the left tail measure substantially exceeds that of the right tail. This asymmetry in the tails under \mathbb{Q} contrasts sharply with the previously discussed approximate symmetric tail behavior under \mathbb{P} . This, of course, is also consistent with the well-known fact that out-of-the-money puts tend to be systematically “overpriced” when assessed by standard pricing models; see, e.g., the discussion in Bates (2009), Bondarenko (2003), Broadie et al. (2009b), and Foresi and Wu (2005). Moreover, all of the three risk-neutral variation measures clearly vary over time, reaching unprecedented levels in Fall 2008. We will study these dynamic patterns more closely below.

5 Tails, Fears and Risk Premia

We begin our discussion of the equity and variance risk premia and the relationship between the \mathbb{P} and \mathbb{Q} measures in Table 1 with a direct comparison of the average intensity for “large” jumps under the two different measures. In an effort to focus on “normal” times and prevent our estimates of the risk premia and investor fears to be unduly influenced by the recent financial market crises, we explicitly exclude the July 2007 through December 2008 part of

²⁸The options always expire on a Saturday with the settlement at the opening of the preceding Friday. The median time-to-expiration of the options in our sample is 14 working days.

²⁹These are taken directly from OptionMetrics and are based on the corresponding mid quotes. In the absence of any deeper out-of-the-money option than the desired moneyness k , we use the implied volatility for the deepest out-of-the-money option that is available as a proxy.




Figure 1: *Implied Measures*. The estimates are based on S&P 500 options spanning the period from January 1996 to December 2008, for a total of 3,237 trading days. The log-moneyness of the options used for the left and right tail measures are fixed at $k = 0.9$ and $k = 1.1$, respectively.

the sample from this estimation.

As is immediately evident from Table 1, the estimated frequencies for the occurrences of “large” jumps are orders of magnitude smaller than the implied jump intensities under the risk-neutral measure.³⁰ Since we explicitly rely on the frequently occurring “medium-sized” jumps and EVT for meaningfully uncovering the \mathbb{P} jump tails, this marked difference cannot simply be attributed to a standard Peso type problem and a non-representative sample. The table also reveals a strong sense of investor fears, as manifest by the much larger intensities for the left jump tails under \mathbb{Q} and the willingness to protect against negative jumps. For instance, the estimated risk-neutral jump intensity for negative jumps less than -20% is more than twelve times larger than that for positive jumps in excess of 20% , and the difference is highly statistically significant. The differences in the left and right \mathbb{P} jump tails, on the other hand, are not nearly as dramatic and also statistically insignificant.

Table 1: Mean Jump Intensities

The jump intensities are reported in annualized units. The jump sizes are in percentage changes in the price level. The estimates under \mathbb{Q} are based on S&P 500 options data from January 1996 through June 2007, while the estimates under \mathbb{P} rely on high-frequency 5-minute S&P futures prices from January 1990 through June 2007. Standard errors for the estimates are reported in parentheses.

| Jump Size | Under \mathbb{Q} | Under \mathbb{P} |
|------------|--------------------|--------------------|
| $> 7.5\%$ | 0.5551 (0.0443) | 0.0098 (0.0142) |
| $> 10\%$ | 0.2026 (0.0206) | 0.0050 (0.0083) |
| $> 20\%$ | 0.0069 (0.0014) | 0.0010 (0.0022) |
| $< -7.5\%$ | 0.9888 (0.0525) | 0.0036 (0.0048) |
| $< -10\%$ | 0.5640 (0.0346) | 0.0017 (0.0026) |
| $< -20\%$ | 0.0862 (0.0084) | 0.0002 (0.0005) |

Although the option implied risk-neutral intensities for positive jumps are much smaller than the intensities for negative jumps, they are still much higher than the estimated intensities for positive jumps under the statistical measure. This contrasts with previous estimates from parametric models, which typically imply trivial and even negative premia for positive jumps. Most of these findings, however, have been based on tightly parameterized models that constrain the overall jump intensity, i.e., the intensity for jumps of any size, to be the same under the risk-neutral and statistical measures. Thus, if negative jumps are priced, in the sense that their risk-neutral intensity exceeds their statistical counterpart, this automatically implies that positive jumps will carry no or even a negative risk premium. Meanwhile, our non-parametric approach, which doesn’t restrict the shape and relation between the \mathbb{P} and \mathbb{Q} jumps, suggest that the large positive jumps do in fact carry a premia, albeit of a

³⁰This is consistent with the aforementioned earlier empirical evidence reported in Ait-Sahalia et al. (2001); see their Section 3.5.

much smaller magnitude than the negative jumps.³¹




Figure 2: *Equity and Variance Risk Premia due to “Large” Jumps*. The estimates for the risk premia are based on 5-minute S&P 500 futures prices and options. The values for the premia in the shaded area, corresponding to July 2007 through December 2008, are based on the parameter estimates for the \mathbb{P} and \mathbb{Q} measures obtained with data through June 2007. The log-moneyness of the options used for the left and right tails are fixed at $k = 0.9$ and $k = 1.1$, respectively.

Turning next to the actual risk premia, we plot in Figure 2 the components of the equity and variance risk premia that are due to compensation for rare events; i.e., $ERP(k)$ and $VRP(k)$, respectively. As mentioned before, we explicitly exclude the last eighteen months of the sample in the estimation of the volatility forecasting model and different jump tail measures, so that the premia in the shaded parts of the figure, coinciding with the advent of the recent financial market crises, are based on the “in-sample” estimates with data through June 2007 only. Also, for ease of interpretation, both of the premia are reported at a monthly

³¹As discussed further below, the temporal variation in the jump intensities means that positive jumps may indeed carry a risk premium. This is reminiscent of the U-shape pattern in the projection of the pricing kernel on the space of market returns previously documented in Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002), among others.

frequency based on 22-day moving averages of the corresponding daily estimates.³²

There are obvious similarities in the general dynamic dependencies in the two jump risk premia, with many of the peaks in $ERP(k)$ coinciding with the troughs in $VRP(k)$. There are also some important differences, however, in the way in which the compensation for the tail events manifest themselves in the two premia. We begin with a more detailed discussion of the equity jump risk premium depicted in the top panel.

5.1 Equity Jump Risk Premium

Most of the peaks in the equity jump risk premium are readily associated with specific economic events. In particular, the October 1997 “mini-crash” and the August-September 1998 turmoil associated with the Russian default and the LTCM debacle both resulted in sharp, but relatively short-lived, increases in the required compensation for tail risks. A similar, albeit much smaller, increase is observed in connection with the September 11, 2001, attacks. Interestingly, the magnitude of the equity jump risk premium in October 2008 is about the same as the premium observed in August-September 1998.

To better understand where this compensation for tail risk is coming from, it is instructive to decompose the total jump risk premium into the parts associated with negative and positive jumps, say $ERP_t^+(k)$ and $ERP_t^-(k)$, respectively. This decomposition, depicted in Figure 3, shows that although the behavior of the two tails are clearly related, the contributions to the overall risk premium are far from symmetric. Most noticeably, the October 1997 “mini-crash” seems to be associated with an increased fear among investors of additional sharp market declines and a corresponding peak in $ERP_t^-(k)$, while there is hardly any effect on $ERP_t^+(k)$. Similarly, the August-September 1998 market turmoil had a much bigger impact on $ERP_t^-(k)$ than it did on $ERP_t^+(k)$. Conversely, the dramatic stock market declines observed in July and September 2002 in connection with the burst of the dot-com “bubble” resulted in peaks and troughs in both $ERP_t^+(k)$ and $ERP_t^-(k)$, with less of a discernible impact on the total jump risk premium $ERP_t(k)$. Likewise, the March 2003 start of the Second Gulf War is hardly visible in $ERP_t(k)$, but clearly influenced both $ERP_t^+(k)$ and $ERP_t^-(k)$. Most dramatic, however, the sharp increase in $ERP_t(k)$ associated with the Fall 2008 financial crises masks even larger offsetting changes in $ERP_t^+(k)$ and $ERP_t^-(k)$ that dwarf the separate jump risk premia observed during the rest of the sample. As such, this directly underscores the notion that investors simply didn’t know what the “right” price was at the time.

³²The time-to-expiration of the options used in the calculation of the risk premia changes systematically over the month, inducing a monthly periodicity in the day-by-day estimates. To illustrate, consider the popular jump-diffusion model of Duffie et al. (2000). For simplicity, suppose there is only a single continuous volatility factor. Let $\kappa^{\mathbb{P}}$ and $\theta^{\mathbb{P}}$ denote the mean-reversion parameter and the mean of σ_t^2 under \mathbb{P} , with the corresponding parameters under \mathbb{Q} denoted by $\kappa^{\mathbb{Q}}$ and $\theta^{\mathbb{Q}}$, respectively. The unconditional mean of the risk premium may then be expressed as $K_0 + K_1(\theta^{\mathbb{P}} - \theta^{\mathbb{Q}})(1 - e^{-\kappa^{\mathbb{Q}}(T-t)})(\kappa^{\mathbb{Q}}(T-t))^{-1}$, where the two constants K_0 and K_1 are determined by the risk-premia parameters. This expected premium obviously depends on the horizon and the relative import of the continuous stochastic volatility component, in turn inducing a monthly periodicity in the daily measures.

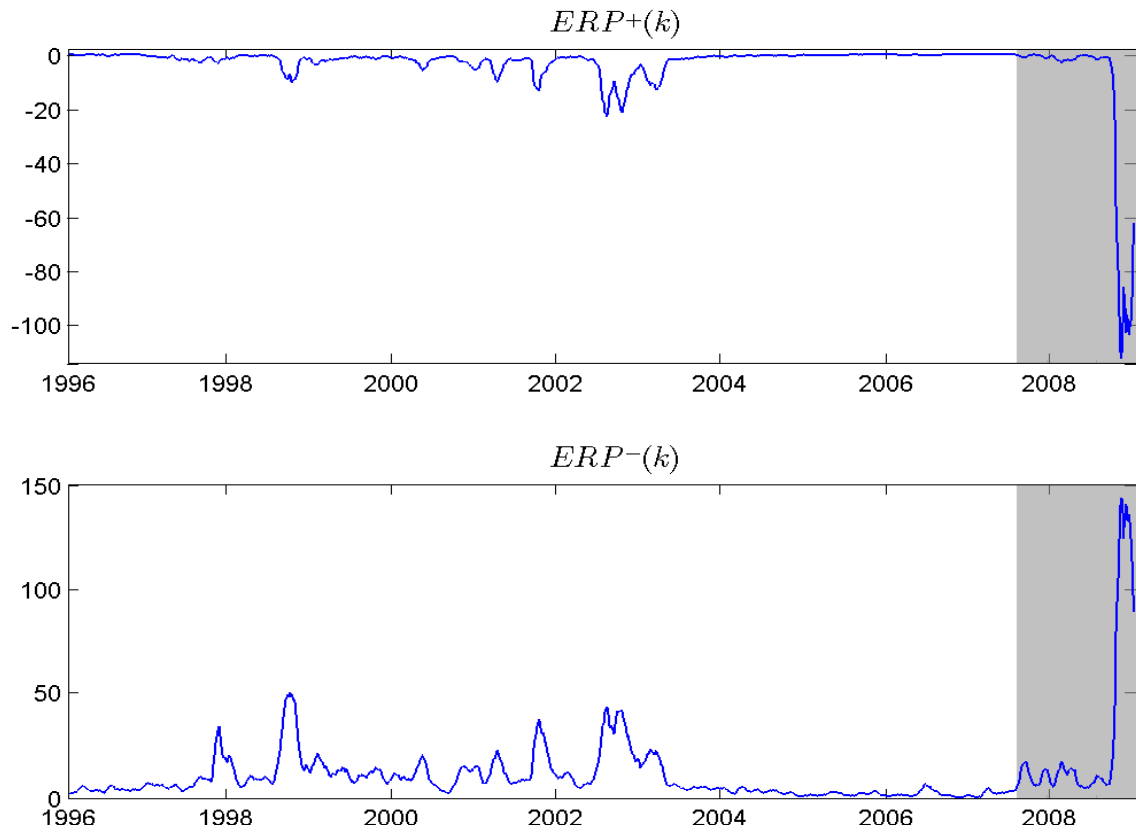


Figure 3: *Decomposition of Equity Jump Premia.* The estimates for the equity risk premia components are based on 5-minute S&P 500 futures prices and options. The values for the premia in the shaded area, corresponding to July 2007 through December 2008, are based on the parameter estimates for the \mathbb{P} and \mathbb{Q} measures obtained with data through June 2007. The log-moneyness of the options used for the left and right tails are fixed at $k = 0.9$ and $k = 1.1$, respectively.

Looking at the typical sample values, again excluding the July 2007-December 2008 part of the sample, the median of the estimated equity risk premia due to rare events equals 5.2%.³³ This is quite high, and compared to the prototypical estimate of 8% for the equity risk premia in postwar U.S. data (see, e.g. Cochrane (2005)), our results imply that fears of rare events account for roughly two-thirds of the total expected excess return. These numbers contrast with the estimation results reported in Broadie et al. (2009a), which imply a mean $ERP_t(k)$ of 1.85%, and the estimates in Eraker (2004), which imply a premium for rare events of only 0.65%.³⁴

At a more general level, our findings suggest that excluding fears and tail events, the magnitude of the equity risk premium is quite compatible with the implications from standard consumption-based asset pricing models with “reasonable” levels of risk aversion. Of course, this raises the question of why investors price tail risk so high? Recent pricing models that pay special attention to tail risks include Pan et al. (2005) and Bates (2008), the work by Bansal and Shaliastovich (2009) and Sizova (2009) based on the concept of confidence risk, and Drechsler (2011) who emphasizes the role of time-varying model uncertainty in amplifying the impact of jumps.

Further corroborating these ideas, we show next that fears of rare events account for an even larger fraction of the historically “large” and difficult to explain variance risk premium.

5.2 Variance Jump Risk Premium

The general dynamic dependencies in the jump variance risk premium depicted in the second panel in Figure 2 fairly closely mirror those of the jump equity premium in the first panel. Similarly, decomposing $VRP_t(k)$ into the parts associated with negative and positive jumps, say $VRP_t^+(k)$ and $VRP_t^-(k)$, respectively, reveal a rather close coherence between the two tail measures.³⁵ With the notable exception of the aforementioned October 1997 “mini-crash” and the August-September 1998 Russian default and LTCM debacle, most of the spikes coincide between the two time series plotted in Figure 4. The magnitude of $VRP_t^-(k)$, of course, typically exceeds that of $VRP_t^+(k)$ by a factor of 2-4.³⁶ Do these findings of large tail risk premia imply a special treatment, or compensation, for rare events? If “large”, i.e., jump-related, risks were treated the same as “small”, i.e., diffusive, risks, what would be the magnitudes of the tail risk premia?

To begin answering these questions, it is instructive to further decompose the variance

³³Including the more recent period the median increases to 5.6%.

³⁴Both of these numbers are based on a specific stochastic volatility model with correlated jumps, referred to in the literature as a SVCJ type model, in which the parametric form of the jumps effectively implies an exponential decay.

³⁵This isn’t simply an artefact of the assumption for the \mathbb{P} jump intensity in equation (3.6). The \mathbb{P} tail moments are orders of magnitude smaller from the moments under \mathbb{Q} , and as such the dynamic dependencies seen in the figure are largely driven by the dependencies in the \mathbb{Q} jump intensity estimated under the less restrictive assumption in (2.6).

³⁶This is also in line with the recent empirical findings in Andersen and Bondarenko (2010) related to their so-called up and down implied variance measures.

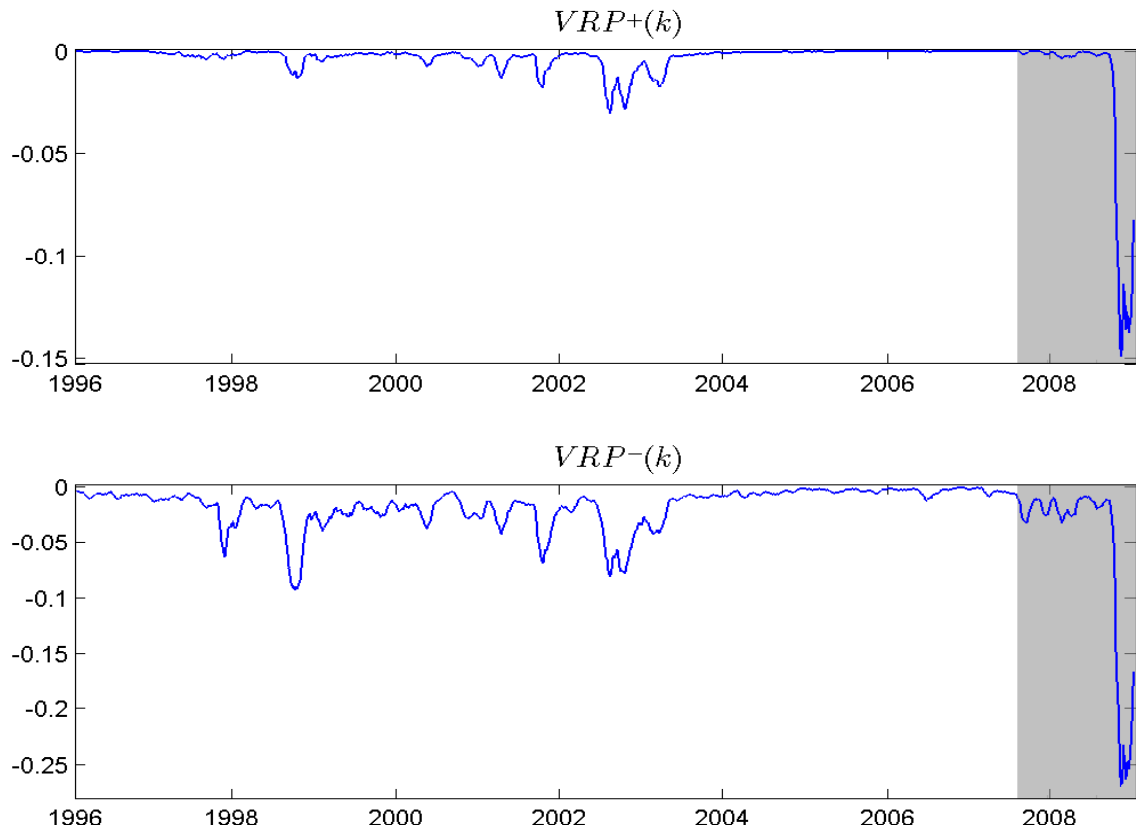


Figure 4: *Decomposition of Variance Jump Premia.* The estimates for the variance risk premia components are based on 5-minute S&P 500 futures prices and options. The values for the premia in the shaded area, corresponding to July 2007 through December 2008, are based on the parameter estimates for the \mathbb{P} and \mathbb{Q} measures obtained with data through June 2007. The log-moneyness of the options used for the left and right tails are fixed at $k = 0.9$ and $k = 1.1$, respectively.

jump tail premia. In particular, from the formal definition of $VRP_t^+(k)$, it follows that

$$\begin{aligned} VRP_t^+(k) &= \frac{1}{T-t} \left(\mathbb{E}_t^{\mathbb{P}} \int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{P}}(dx) - \mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{P}}(dx) \right) \\ &\quad + \frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{P}}(dx) - \int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{Q}}(dx) \right), \end{aligned} \quad (5.1)$$

with a similar decomposition available for $VRP_t^-(k)$.

The first term on the right-hand-side in equation (5.1) reflects the compensation for time-varying jump intensity risk. It mirrors the compensation demanded by investors for the continuous part of the quadratic variation, i.e., $VRP_t^c(k)$ in equation (1.10). As such it is naturally associated with investor's willingness to hedge against changes in the investment opportunity set. Specifically, for the jump intensity process in equation (3.6), the relevant difference between the \mathbb{P} and \mathbb{Q} expectations takes the form,

$$\mathbb{E}_t^{\mathbb{P}} \int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{P}}(dx) - \mathbb{E}_t^{\mathbb{Q}} \int_t^T \int_{x>k} x^2 ds \nu_s^{\mathbb{P}}(dx) = \alpha_1^+ \int_{x>k} x^2 \nu^{\mathbb{P}}(x) dx \left(\mathbb{E}_t^{\mathbb{P}} \int_t^T \sigma_s^2 ds - \mathbb{E}_t^{\mathbb{Q}} \int_t^T \sigma_s^2 ds \right). \quad (5.2)$$

For Lévy type jumps, i.e, jumps with constant jump intensity, this part of the variance jump risk premium will be identically equal to zero.

By contrast, the second term on the right-hand-side in equation (5.1) involves the wedge between the risk-neutral and objective jump intensities. But, this difference is evaluated under the same probability measure, and as such it is effectively purged of the premia due to temporal variation in the jump intensities. If jumps were treated the same as diffusive risks, this component would be zero, since the predictable quadratic variation of the continuous price does not change when changing the measure from \mathbb{P} to \mathbb{Q} .

In more technical terms, the quadratic variation of the log-price process may be split into its predictable component and a residual martingale component solely due to jumps,

$$QV_{[t,T]} = \int_t^T \sigma_s^2 ds + \int_t^T \int_{\mathbb{R}} x^2 ds \nu_s^{\mathbb{P}}(dx) + \int_t^T \int_{\mathbb{R}} x^2 \tilde{\mu}(ds, dx) \equiv \langle f_s, f_s \rangle_{[t,T]} + \int_t^T \int_{\mathbb{R}} x^2 \tilde{\mu}(ds, dx),$$

where $\langle f_s, f_s \rangle_{[t,T]}$ denotes the predictable quadratic variation; see, e.g., Jacod and Shiryaev (2003). The pricing of the first term, $\langle f_s, f_s \rangle_{[t,T]}$, is naturally associated with changes in the investment opportunity set, while the pricing of the martingale part reflect any "special" treatment of jump risk.

The high-frequency based estimation results discussed in the previous section indicate significant temporal variation in the tail jump intensities. Part of the tail risk premia will therefore be coming from the first term in equation (5.1). However, our estimation results, see Table 6 in the Appendix, also suggest that the left and right tails behave quite similarly




Figure 5: *Investor Fears*. The estimates for the $FI(k)$ investor fear index are based on 5-minute S&P 500 futures prices and options. The values in the shaded area, corresponding to July 2007 through December 2008, are based on the parameter estimates for the \mathbb{P} and \mathbb{Q} measures obtained with data through June 2007. The log-moneyness of the options used for the left and right tails are fixed at $k = 0.9$ and $k = 1.1$, respectively.

under \mathbb{P} , so that in particular³⁷

$$\alpha_1^+ \int_{x>k} x^2 \nu^{\mathbb{P}}(x) dx \approx \alpha_1^- \int_{x<-k} x^2 \nu^{\mathbb{P}}(x) dx.$$

Combining this approximation with (5.2) and the comparable expression for the left tail measure, it follows that the difference

$$FI_t(k) = VRP_t^-(k) - VRP_t^+(k), \quad (5.3)$$

will be largely void of the risk premia due to temporal variation in the jump intensities. Consequently, $FI_t(k)$ may be interpreted as a direct measure of investor fears, or “crash-o-phobia.”³⁸

The corresponding plot in Figure 5 shows that the sharpest increase in investor fears over the sample did indeed occur during the recent financial crises. However, the Russian default and LTCM collapse in August-September 1998 resulted in a spike in the fear index of almost 2/3 of its recent peak. Slightly less dramatic increases are manifest in connection with the October 1997 “mini-crash”, September 11, 2001, and the summer of 2002 and the burst of the dot-com “bubble”. The figure also reveals systematically low investor fears from 2003 through mid-2007, corresponding to the general run-up in market values.

The average sample values reported in Table 2 further corroborate the idea that much of the variance risk premium comes from “crash-o-phobia”, or special compensation for tail risks. Excluding the post July 2007 crises period, the average sample variance risk premium equals $\mathbb{E}(VRP_t) = 0.020$. Of that total premium, $\mathbb{E}(VRP_t^+(k))/\mathbb{E}(VRP_t) = 14.8\%$ is due to compensation for right tail risks, while $\mathbb{E}(VRP_t^-(k))/\mathbb{E}(VRP_t) = 88.4\%$ comes from left tail risks. Looking at the difference, our results therefore suggest that close to 3/4 of the variance risk premium may be attributed to investor fears, as opposed to more standard rational based pricing arguments.³⁹

Comparing these numbers to the implications from the aforementioned parametric model estimates reported in the literature, the study by Broadie et al. (2009a) imply that $\mathbb{E}(VRP_t(k))/\mathbb{E}(VRP_t)$

³⁷Note that this result is not “by construction” as assumption (3.6) allows for asymmetric tails since the left and the right tails are allowed to have different loadings on the constant and time-varying part of the jump intensity.

³⁸Rubinstein (1994) first attributed the smirk like pattern in post October 1987 Black-Scholes implied volatilities for the aggregate market portfolio when plotted against the degree of moneyness as evidence of “crash-o-phobia”; see also Foresi and Wu (2005) for more recent related international empirical evidence.

³⁹Note that these proportions are maturity specific, and based on a median maturity of 14 days. To check the sensitivity of our results with respect to the maturity and the statistical uncertainty associated with our non-parametric inference procedure, we redid the estimation using monthly data comprised of options on the fifth business day of each month. The resulting monthly estimates for $\mathbb{E}(VRP_t^+(k))/\mathbb{E}(VRP_t) = 8.0\%$ and $\mathbb{E}(VRP_t^-(k))/\mathbb{E}(VRP_t) = 60.8\%$ still imply that the fear component accounts for more than half of the variance risk premium. Thus, even though the relative importance of $VRP_t(k)$ drops somewhat, the general findings remain intact. We note that these differences are also consistent with the biases observed in the Monte Carlo reported in Section 6.

Table 2: Variance Risk Decomposition

The average sample estimates are reported in annualized form, with standard errors in parentheses. The estimates under \mathbb{Q} are based on S&P 500 options data from January 1996 through June 2007, while the estimates under \mathbb{P} rely on 5-minute high-frequency S&P 500 futures data from January 1990 through June 2007.

| Under \mathbb{Q} | |
|---|---------------------|
| $\mathbb{E} \left(\frac{1}{T-t} \int_t^T \int_{x > \ln 1.1} x^2 \mathbb{E}_t^{\mathbb{Q}} \nu_s^{\mathbb{Q}}(dx) ds \right)$ | 0.0032 (0.0004) |
| $\mathbb{E} \left(\frac{1}{T-t} \int_t^T \int_{x < \ln 0.9} x^2 \mathbb{E}_t^{\mathbb{Q}} \nu_s^{\mathbb{Q}}(dx) ds \right)$ | 0.0180 (0.0015) |
| $\mathbb{E}(QV_t^{\mathbb{Q}})$ | 0.0472 (0.0046) |
| Under \mathbb{P} | |
| $\mathbb{E} \left(\frac{1}{T-t} \int_t^T \int_{x > \ln 1.1} x^2 \mathbb{E}_t^{\mathbb{P}} \nu_s^{\mathbb{P}}(dx) ds \right)$ | 0.000196 (0.000483) |
| $\mathbb{E} \left(\frac{1}{T-t} \int_t^T \int_{x < \ln 0.9} x^2 \mathbb{E}_t^{\mathbb{P}} \nu_s^{\mathbb{P}}(dx) ds \right)$ | 0.000062 (0.000115) |
| $\mathbb{E}(QV_t^{\mathbb{P}})$ | 0.0269 (0.0031) |

24.4%, while the results in Eraker (2004) suggest that $\mathbb{E}(VRP_t(k))/\mathbb{E}(VRP_t) = 20.0\%$.⁴⁰ These numbers obviously differ quite dramatically from the results obtained with our non-parametric approach, which effectively suggest that all, or most, of the variance risk premium may be attributed to the compensation for jump tail risk. Importantly, however, existing parametric studies have generally constrained the wedge between the risk-neutral and statistical jump distributions to a negative shift in the mean jump size;⁴¹ see, e.g., the discussion in Singleton (2006). The results reported here indicate that the difference in the jump behavior under the two probability measures is much more intricate than that, and that time-varying jump intensities play an especially important role for properly understanding the risk premia and the inherent fear component.

We next check the robustness of our findings with respect to the general assumptions concerning the \mathbb{P} and \mathbb{Q} dynamic tail dependencies, and the specific asymptotic approximations underlying our jump tail estimates.

6 Robustness Checks

Our approach for estimating the tail risk premia and the investor fears index is extremely flexible, the only parametric assumptions relating to the very general form of the time-varying

⁴⁰These numbers are calculated on the basis of the specific parameter estimates for the SVCJ models reported in the two studies, with the horizon fixed at 14 days, corresponding to the median time-to-expiration of the options used in our analysis, along with $k = 10\%$.

⁴¹Broadie et al. (2009a) also consider changing the variance of the jumps.

jump intensities for the \mathbb{P} and \mathbb{Q} jump tails in equations (3.6) and (2.6), respectively. We begin our discussion below with a series of formal rank tests for the empirical validity of these assumptions that underly our estimation of the RT_t and LT_t tail measures. We then present the summary results from a detailed Monte Carlo simulation study designed to investigate the accuracy of our \mathbb{Q} jump tail estimation. We also briefly comment on some additional robustness checks related to the use of lower frequency daily data and shorter-lived options.

6.1 Time-Variation in \mathbb{P} Jump Tails

The relevant assumption in equation (3.6) restricts the intensities for the positive and negative jumps to be affine functions, possibly with different coefficients, of the latent instantaneous volatility σ_t^2 . This in turn implies that the left and right realized jump variances, $\int_{t+\pi_t}^{t+1} \int_{x>0} x^2 \mu(ds, dx)$ and $\int_{t+\pi_t}^{t+1} \int_{x<0} x^2 \mu(ds, dx)$, only depend upon the past through the temporal variation in the integrated stochastic volatility. Consequently the unconditional (auto)covariances equal,

$$\text{Cov} \left(\int_{t+\pi_t}^{t+1} \int_{x>0} x^2 \mu(ds, dx), \int_{t-1+\pi_t}^t \sigma_s^2 ds \right) = K_1 \text{Cov} \left(\int_{t+\pi_t}^{t+1} \sigma_s^2 ds, \int_{t-1+\pi_t}^t \sigma_s^2 ds \right),$$

and,

$$\begin{aligned} \text{Cov} \left(\int_{t+\pi_t}^{t+1} \int_{x>0} x^2 \mu(ds, dx), \int_{t-1+\pi_t}^t \int_{x>0} x^2 \mu(ds, dx) \right) = \\ K_1 \text{Cov} \left(\int_{t+\pi_t}^{t+1} \sigma_s^2 ds, \int_{t-1+\pi_t}^t \int_{x>0} x^2 \mu(ds, dx) \right), \end{aligned}$$

where $K_1 = (\alpha_1^+)^2 \int_{x>0} x^2 \nu^{\mathbb{P}}(x) dx$. Now, defining the vector

$$\mathbf{Z}_t \equiv \left(\int_{t+\pi_t}^{t+1} \sigma_s^2 ds, \int_{t+\pi_t}^{t+1} \int_{x>0} x^2 \mu(ds, dx), \int_{t+\pi_t}^{t+1} \int_{x<0} x^2 \mu(ds, dx) \right)',$$

it follows that the autocovariance matrix for \mathbf{Z}_t should be of reduced rank.⁴² We state this testable implication in the following proposition.

Proposition 2 *If σ_t^2 is a covariance stationary and ergodic process, and $\int_{\mathbb{R}} x^2 \nu^{\mathbb{P}}(x) dx < \infty$, the assumption in equation (3.6) implies that $\text{rank}(\text{Cov}(\mathbf{Z}_t, \mathbf{Z}_{t-1})) = 1$.*

In practice, of course, the vector Z_t is not directly observable. However, as discussed in Section 3 above, it may be accurately estimated by (CV_t, RJV_t, LJV_t) . Even so, this vector will invariably include some estimation error, which will render the resulting autocovariance matrix of full rank even if $\text{Cov}(\mathbf{Z}_t, \mathbf{Z}_{t-1})$ is not. Hence, in order to formally test whether the deviations from reduced rank in the actually observed sample autocovariance matrices for (CV_t, RJV_t, LJV_t) , and subsets thereof, are statistically significant, we rely on the test

⁴²The possibility of jumps in σ_t^2 means that the covariance matrix of Z_t will not necessarily be of reduced rank.

developed by Gill and Lewbel (1992); further details concerning the implementation are given in Appendix D.⁴³

Table 3: Rank Tests for Affine \mathbb{P} Jump Intensities

The table reports the results of the Gill and Lewbel (1992) rank test, as detailed in Appendix D. The realized variation measures are based on 5-minute high-frequency S&P 500 futures data from January 1990 through December 2008. The test has a limiting χ_d^2 distribution under the null hypothesis that the autocovariance matrix has reduced rank one. The calculation of the asymptotic variance for the autocovariance matrices are based on a Parzen kernel with a lag-length of 70.

| Autocovariance Matrix | Test Stat. | d | p-value |
|------------------------|------------|---|---------|
| (CV_t, RJV_t, LJV_t) | 3.7937 | 2 | 0.1500 |
| (CV_t, RJV_t) | 1.2616 | 1 | 0.2613 |
| (CV_t, LJV_t) | 0.2590 | 1 | 0.6108 |
| (RJV_t, LJV_t) | 3.5299 | 1 | 0.0603 |

None of the resulting test statistics reported in Table 3 reject the reduced rank hypothesis. As such, the results do not provide any significant evidence against the affine restriction in equation (3.6) related to the actual jump intensities.⁴⁴

6.2 Time-Variation in \mathbb{Q} Jump Tails

In contrast to the affine dependencies for the \mathbb{P} jump intensities underlying our estimates of the statistical tail measures, our estimates for the risk-neutral tail measures rely on the less restrictive condition in equation (2.6). This assumption merely requires the temporal variation to be the same across different jump sizes, and as such is trivially satisfied by all parametric models hitherto estimated in the literature. This level of generality for the \mathbb{Q} jump measures may be particularly important to properly accommodate the seemingly complex dynamic dependencies in investor fears.

In order to more formally investigate this hypothesis, suppose that the \mathbb{Q} jump intensities satisfied the analogous affine dependencies assumed for the \mathbb{P} jump measure,

$$\nu_t^{\mathbb{Q}}(dx) = (\alpha_0^{+\mathbb{Q}}1_{\{x>0\}} + \alpha_0^{-\mathbb{Q}}1_{\{x<0\}} + (\alpha_1^{+\mathbb{Q}}1_{\{x>0\}} + \alpha_1^{-\mathbb{Q}}1_{\{x<0\}})\sigma_t^2) \nu^{\mathbb{Q}}(x)dx, \quad (6.1)$$

⁴³This test is relatively easy to implement, and it overcomes some of the computational difficulties associated with other tests for rank deficiency based on the smallest eigenvalues of the matrix. The Monte Carlo simulation evidence discussed in Gill and Lewbel (1992) also suggests that the test exhibits desirable size and power properties.

⁴⁴The test statistics reported in the table are based on the realized variation measures over the full sample. However, very similar test statistics are obtained for the shorter January 1990 through June 2007 sample employed in the tail estimation. We also implemented the same set of tests with all of the variables replaced by their square-roots; i.e., $(CV_t^{1/2}, RJV_t^{1/2}, LJV_t^{1/2})$. All of these tests strongly rejected the reduced rank hypothesis, highlighting the power of the procedure.

where $\nu^{\mathbb{Q}}(x)$ denotes a valid Lévy density, and $\alpha_0^{\pm\mathbb{Q}}$ and $\alpha_1^{\pm\mathbb{Q}}$ are nonnegative constants.⁴⁵ In parallel to the definition of Z_t in Proposition 2 above, define the vector of risk-neutral measures, $\mathbf{Y}_t \equiv (RT_t^{\mathbb{Q}}(k), LT_t^{\mathbb{Q}}(k), QV_t^{\mathbb{Q}})'$. If (6.1) holds true, $RT_t^{\mathbb{Q}}(k)$, $LT_t^{\mathbb{Q}}(k)$ and $QV_t^{\mathbb{Q}}$ should all be linear functions of the risk-neutral expected integrated volatility $\mathbb{E}_t^{\mathbb{Q}}\left(\int_t^T \sigma_s^2 ds\right)$. Consequently, the covariance matrix, and equivalently the lag-one autocovariance matrix, for \mathbf{Y}_t should be of reduced rank one. If (6.1) doesn't hold, $QV_t^{\mathbb{Q}}$ may still depend linearly on the integrated volatility $\mathbb{E}_t^{\mathbb{Q}}\left(\int_t^T \sigma_s^2 ds\right)$. However, in this situation additional time-varying factors will typically affect \mathbf{Y}_t , so that $\text{rank}(\text{Var}(\mathbf{Y}_t)) \equiv \text{rank}(\text{Cov}(\mathbf{Y}_t, \mathbf{Y}_{t-1})) > 1$. The following testable proposition summarizes this implication.

Proposition 3 *If Y_t is a covariance stationary and ergodic process, and $\int_{|x|>\epsilon} e^x \nu^{\mathbb{Q}}(x) dx < \infty$ for some $\epsilon > 0$, the assumption in equation (6.1) implies that $\text{rank}(\text{Cov}(\mathbf{Y}_t, \mathbf{Y}_{t-1})) = 1$.*

The results from the rank tests for the \mathbb{Q} tail measures are reported in Table 4. As seen from the table, all of the tests clearly reject the affine structure for the risk-neutral jump intensities.⁴⁶ This, of course, contrasts sharply with the comparable test results for the \mathbb{P} jump intensities in Table 3, which generally support the affine structure.

Table 4: Rank Tests for Affine \mathbb{Q} Jump Intensities

The table reports the results of the Gill and Lewbel (1992) rank test, as detailed in Appendix D. The variation measures are based on S&P 500 options data from January 1996 through December 2008. The test has a limiting χ_d^2 distribution under the null hypothesis that the autocovariance matrix has reduced rank one. The calculation of the asymptotic variance for the autocovariance matrices are based on a Parzen kernel with a lag-length of 70.

| Autocovariance Matrix | Test Stat. | d | p-value |
|---|------------|---|---------|
| $(QV_t^{\mathbb{Q}}, RT_t^{\mathbb{Q}}, LT_t^{\mathbb{Q}})$ | 65.8861 | 2 | 0.0000 |
| $(QV_t^{\mathbb{Q}}, RT_t^{\mathbb{Q}})$ | 29.3789 | 1 | 0.0000 |
| $(QV_t^{\mathbb{Q}}, LT_t^{\mathbb{Q}})$ | 31.7486 | 1 | 0.0000 |
| $(RT_t^{\mathbb{Q}}, LT_t^{\mathbb{Q}})$ | 12.4262 | 1 | 0.0000 |

Taken as a whole, the robustness tests in Tables 3 and 4 pertaining to the temporal variation in the jump intensities are all consistent with the key minimal assumptions underlying our non-parametric inference procedures. Importantly, the test results clearly point to the existence of much more complex dynamic dependencies in the risk-neutral jump tails than

⁴⁵This assumption, together with equation (3.6), imply that the temporal variation in all risk premia would be spanned by the level of the stochastic volatility.

⁴⁶The tests reported in the table are based on the full sample including date through the end of 2008. As before, very similar results are obtained when excluding the July 2007 through December 2008 part of the sample.

hitherto entertained in the literature, and in the form of dependencies that cannot simply be spanned by the level of the stochastic volatility. As such, this also calls into question the validity of parametric procedures that restrict the form of the tails to be the same across the \mathbb{P} and \mathbb{Q} distributions when trying to understand notions of investor fears, or “crash-o-phobia”, embedded in market prices.

6.3 Estimation of \mathbb{Q} Tail Measures

To gauge the magnitude of the approximation errors in Proposition 1 and the accuracy of the estimates for the $LT_t^{\mathbb{Q}}$ and $RT_t^{\mathbb{Q}}$ tail measures, we simulated artificial option prices from four different parametric models. Model A refers to the classical Merton model with constant volatility and normally distributed compound Poisson jumps. Following Carr and Wu (2003), we fix the continuous volatility at 14%, the jump intensity at 2, the mean log jump size at -5% , and the volatility of the log jump size at 13%, all in annualized units.⁴⁷

The remaining three Models B, C and D all assume that the instantaneous volatility follows a square-root diffusion process,

$$d\sigma_t^2 = \kappa_i(\theta_i - \sigma_t^2)dt + \eta\sigma_t dW_t, \quad i = p, q, \quad (6.2)$$

where $\eta = 0.1$, and the rest of the parameters under the \mathbb{P} and \mathbb{Q} probability measures are fixed at $\theta_p = 0.13^2$, $\kappa_p = 5.04$, and $\theta_q = 0.14^2$, $\kappa_q = 4.3457$, respectively. Following Bates (2000), Model B further assumes normally distributed jumps with jump intensity proportional to the instantaneous variance,

$$\nu_t(dx) = \lambda_i\sigma_t^2 \frac{e^{-0.5(x-\mu_i)^2/\eta_i^2}}{\sqrt{2\pi\eta_i^2}}, \quad i = p, q, \quad (6.3)$$

with the \mathbb{P} and \mathbb{Q} parameters set at $\lambda_p = 2/\theta_p$, $\mu_p = 0$, $\eta_p = 0.0561$, and $\lambda_q = 2/\theta_q$, $\mu_q = -0.05$, and $\eta_q = 0.13$, respectively. These values for the \mathbb{Q} jump parameters ensure that conditional on the mean level of the volatility, the jumps have the same distribution and intensity as in Model A.

Finally, Models C and D postulate a generalized tempered stable process for the jump compensator, as in, e.g., Bates (2009) and Carr et al. (2003),

$$\nu_t(dx) = \frac{\sigma_t^2}{|x|^{\alpha+1}} (c^- e^{-M_i|x|} 1_{\{x<0\}} + c^+ e^{-N_i x} 1_{\{x>0\}}), \quad i = p, q, \quad (6.4)$$

where for Model C the parameters are fixed at $M_p = N_p = 13.0543$, $M_q = 3$, $N_q = 8.1313$, $\alpha = 0.0$, and $c^- = c^+ = 0.3600/\theta_p$, while for Model D we set $M_p = N_p = 19.5026$, $M_q = 3$, $N_q = 9.1990$, $\alpha = 0.5$, and $c^- = c^+ = 0.2053/\theta_p$. The \mathbb{Q} jump parameters are chosen so that the expected values of $LT_t^{\mathbb{Q}}$ and $RT_t^{\mathbb{Q}}$ implied by $\nu_t^{\mathbb{Q}}(dx)$ are in line with their corresponding sample means from our data. The \mathbb{P} jump parameters imply that the mean value of the

⁴⁷The only difference from Carr and Wu (2003) is the reduction in the mean log-jump size from -10% to -5% , in order to more closely resemble the sample averages in our data.

daily P jump variation equals 0.25 in daily percentage units, which is generally consistent with recent non-parametric empirical evidence reported in the literature.

Turning to the results, Table 5 reports the Median and Median Absolute Deviation (MAD) estimates for the two tail measures obtained across 5,000 simulated “days” for each of the four models for options with $T - t$ set to 15-days to maturity and k equal to 0.9 and 1.1 for the left and right tails, respectively.⁴⁸ These characteristics correspond directly to the average maturity time and moneyness for the options used in the actual estimation. As is evident from the table, the left skew in the risk-neutral jump distributions render the left tail measures 2 to 3.5 times larger than the corresponding right tail measures and generally also easier to estimate. The slight upward bias observed in both tails is directly attributable to the continuous variation invariably “clouding” the identification of the true jump tails. All-in-all, however, the results clearly underscore the accuracy of the theoretical approximations underlying our estimation of the jump tails.⁴⁹

Table 5: Estimation of Risk-Neutral Tail Measures

The table reports summary statistics for the estimation of the LT_t and RT_t measures based on out-of-the-money options computed with simulated data from Models A-D as described in the text. Time-to-maturity $T - t$ is fixed at 15 days, and the log-moneyness k for the left and right tails are set at 0.9 and 1.1, respectively. The length of the simulated series is 5,000 days. MAD denotes the Median Absolute Deviation.

| | $LT_t^{\mathbb{Q}}$ | | | $RT_t^{\mathbb{Q}}$ | | |
|---------|---------------------|--------|--------|---------------------|--------|--------|
| | True | Median | MAD | True | Median | MAD |
| Model A | 0.0484 | 0.0496 | — | 0.0201 | 0.0238 | — |
| Model B | 0.0418 | 0.0437 | 0.0072 | 0.0173 | 0.0205 | 0.0039 |
| Model C | 0.0407 | 0.0406 | 0.0056 | 0.0112 | 0.0128 | 0.0021 |
| Model D | 0.0380 | 0.0408 | 0.0070 | 0.0101 | 0.0139 | 0.0032 |

6.4 Daily Data and Shorter Options

To check the robustness of our main empirical findings with respect to the use of intraday data for the extraction of the \mathbb{P} jump tails, we redid the estimation of the tail parameters using only daily data. Restricting the analysis to a coarser daily frequency invariably handicaps the detection of jumps and the corresponding jump tail estimation. Nonetheless, the key features pertaining to the behavior of the jump tails remain intact, albeit not as precisely

⁴⁸The results for the Merton Model A are exact, so the MADs are identically equal to zero.

⁴⁹Much more detailed simulation results and further discussion for a range of different maturity times and moneyness are available in the Supplementary Appendix. Not surprisingly, these results reveal that the quality of the tail estimates generally deteriorate the closer to the money and the longer lived are the options used in the estimation.

estimated.⁵⁰

In a similar vein, to corroborate the simulation results discussed above and check the robustness of our empirical results with respect to the options used in extracting the \mathbb{Q} jump tails, we re-estimated the tails with only monthly data. Specifically, by restricting our sample to the 5-th business day of each month, it is possible to reduce the median time-to-maturity of the options from 14 to only 7 business days. Nonetheless, the resulting monthly Investor Fears index looks almost identical to the plot discussed in Section 4 above. Further details concerning these and other robustness checks are again available in the Supplementary Appendix.

7 Conclusion

The recent market turmoil has spurred a renewed interest in the study and economic impact of rare disaster type events. Most of these studies resort to peso type explanations for the occurrence of rare events along with specific parametric modeling assumptions for carrying out the statistical inference. In contrast, the new methodologies developed here are based on flexible non-parametric procedures, relying on high-frequency data and extreme value theory for effectively measuring the actual expected tail events, together with a rich set of options for uncovering the pricing of tail events.

Allowing the data to “speak for themselves,” our empirical findings suggest that much of the historically large equity and variance risk premia may be ascribed to compensation for jump tail risk. These types of risks are poorly described by traditional mean-variance type frameworks. Meanwhile, excluding the part of the risk premia directly attributable to tail events, the average equity and variance premia observed in the data are both in line with the general implications from standard equilibrium based asset pricing models with reasonable preference parameters. Our decomposition of the variance risk premia further suggests that investors fears, or “crash-o-phobia”, play a crucial role in explaining the overall large magnitude of the tail risk premia. At the same time, the perception of jump tail risks appear to vary quite dramatically over time, and changes therein may help explain some of the recent sharp movements in market valuations.

⁵⁰By the basic statistical principle that more data is always better (see, e.g., the pertinent discussion in Zhang, Mykland, and Ait-Sahalia, 2005), the high-frequency data, if used correctly, will *always* result in more accurate estimates.

Appendices

A Estimation of Jump Tails

We start with some general remarks that apply for both the \mathbb{P} and \mathbb{Q} distributions, and then subsequently specialize the discussion to the statistical and risk-neutral jump tail estimation. For ease of notation, we will suppress the dependence on \mathbb{P} and \mathbb{Q} in the Lévy densities $\nu^{\mathbb{P}}(x)$ and $\nu^{\mathbb{Q}}(x)$ in this initial discussion.

To begin, define $\psi^+(x) = e^x - 1$ for $x \in \mathbb{R}_+$ and $\psi^-(x) = e^{-x}$ for $x \in \mathbb{R}_-$. $\psi^+(x)$ maps \mathbb{R}_+ into \mathbb{R}_+ , and $\psi^-(x)$ maps \mathbb{R}_- into $(1, +\infty)$. Both mappings are one-to-one. Further, denote $\nu_{\psi}^+(y) = \frac{\nu(\ln(y+1))}{y+1}$ for $y \in (0, \infty)$ and $\nu_{\psi}^-(y) = \frac{\nu(-\ln y)}{y}$. Finally, denote the tail jump measures

$$\bar{\nu}_{\psi}^{\pm}(x) = \int_x^{\infty} \nu_{\psi}^{\pm}(u) du, \quad (\text{A.1})$$

with $x > 0$ for $\bar{\nu}_{\psi}^+(x)$, and $x > 1$ for $\bar{\nu}_{\psi}^-(x)$.

Our estimation under both \mathbb{P} and \mathbb{Q} are nonparametric and based on the assumption that the tail Lévy densities $\bar{\nu}_{\psi}^{\pm}$ belong to the domain of attraction of an *extreme value distribution*; see, e.g., the general discussion in Embrechts et al. (2001), along with the recent empirical analysis related to extreme market moves in Bakshi et al. (2009). This is a very wide class of jump measures, and it includes most parametric specifications that have previously been used in the literature. The key implication of the assumption is, see, e.g., Smith (1987), that

$$1 - \frac{\bar{\nu}_{\psi}^{\pm}(u+x)}{\bar{\nu}_{\psi}^{\pm}(x)} \stackrel{\text{appr}}{\approx} G(u; \sigma^{\pm}, \xi^{\pm}), \quad u > 0, x > 0, \quad (\text{A.2})$$

where

$$G(u; \sigma^{\pm}, \xi^{\pm}) = \begin{cases} 1 - (1 + \xi^{\pm} u / \sigma^{\pm})^{-1/\xi^{\pm}}, & \xi^{\pm} \neq 0, \sigma^{\pm} > 0 \\ 1 - e^{-u/\sigma^{\pm}}, & \xi^{\pm} = 0, \sigma^{\pm} > 0 \end{cases} \quad (\text{A.3})$$

$G(u; \sigma^{\pm}, \xi^{\pm})$ is the cdf of a *generalized Pareto distribution*. In what follows we assume that the approximation in (A.2) is exact when x is “large” enough; see the formal proofs in Bollerslev and Todorov (2010) for further details concerning the approximation error and the exact conditions under which the error is indeed asymptotically negligible.

Before we turn to our discussion of the estimation of the parameters for the generalized Pareto distribution approximation to the tail Lévy densities, we present several integrals with respect to the Lévy density $\nu(x)$ that appear in $ERP_t(k)$, $VRP_t(k)$, $RT_t^{\mathbb{P}}(k)$, $LT_t^{\mathbb{P}}(k)$, $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$. In particular, assuming that the approximation in (A.2) holds, it follows that for $k > 0$ and some $tr^+ < e^k - 1$,

$$\int_{x>k} (e^x - 1) \nu(x) dx = (e^k - 1) \bar{\nu}_{\psi}^+(e^k - 1) + \int_{\mathbb{R}} (e^x - e^k)^+ \nu(x) dx,$$

$$\begin{aligned}
\int_{\mathbb{R}} (e^x - e^k)^+ \nu(x) dx &= \int_{(0, \infty)} (x - e^k + 1)^+ \nu_{\psi}^+(x) dx = \int_{e^{k-1}}^{\infty} \bar{\nu}_{\psi}^+(x) dx \\
&= \bar{\nu}_{\psi}^+(tr^+) \int_{e^{k-1}}^{\infty} \left(1 - \left(1 - \frac{\bar{\nu}_{\psi}^+(x)}{\bar{\nu}_{\psi}^+(tr^+)} \right) \right) dx \\
&= \bar{\nu}_{\psi}^+(tr^+) \left(\frac{\sigma^+}{1 - \xi^+} \right) \left(1 + \frac{\xi^+}{\sigma^+} (e^k - 1 - tr^+) \right)^{1-1/\xi^+}, \\
\int_{x>k} (e^x - 1 - x) \nu(x) dx &= \int_{x>e^{k-1}} (x - \ln(x+1)) \nu_{\psi}^+(x) dx \\
&= \bar{\nu}_{\psi}^+(e^k - 1)(e^k - 1 - k) + \bar{\nu}_{\psi}^+(tr^+) \frac{\sigma^+}{1 - \xi^+} \left(1 + \frac{\xi^+}{\sigma^+} (e^k - 1 - tr^+) \right)^{1-1/\xi^+} \\
&\quad - \frac{\bar{\nu}_{\psi}^+(tr^+) (\sigma^+)^{1/\xi^+} \xi^+}{(e^k \xi^+)^{1/\xi^+}} {}_2F_1 \left(\frac{1}{\xi^+}; \frac{1}{\xi^+}; 1 + \frac{1}{\xi^+}; \frac{\xi^+ (tr^+ + 1) - 1}{\frac{\xi^+}{\sigma^+} e^k} \right), \\
\int_{x>k} x^2 \nu(x) dx &= k^2 \bar{\nu}_{\psi}^+(e^k - 1) + 2 \int_{x>e^{k-1}} \frac{\ln(x+1)}{x+1} \bar{\nu}_{\psi}^+(x) dx = k^2 \bar{\nu}_{\psi}^+(e^k - 1) + 2 \bar{\nu}_{\psi}^+(tr^+) \times K_1, \\
K_1 &= e^{-k/\xi^+} \xi^+ \left(\frac{\xi^+}{\sigma^+} \right)^{-1/\xi^+} \left[\xi^+ {}_3F_2 \left(\frac{1}{\xi^+}, \frac{1}{\xi^+}, \frac{1}{\xi^+}; 1 + \frac{1}{\xi^+}, 1 + \frac{1}{\xi^+}; \frac{\xi^+ (tr^+ + 1) - 1}{e^k \frac{\xi^+}{\sigma^+}} \right) \right. \\
&\quad \left. + k {}_2F_1 \left(\frac{1}{\xi^+}, \frac{1}{\xi^+}, 1 + \frac{1}{\xi^+}; \frac{\xi^+ (tr^+ + 1) - 1}{e^k \frac{\xi^+}{\sigma^+}} \right) \right],
\end{aligned}$$

where ${}_3F_2(, , , ; , ;)$ and ${}_2F_1(, , ;)$ are the hypergeometric functions. The above integrals are all well defined provided that $\int_{x>\epsilon} x \nu_{\psi}^+(x) dx < \infty$ for some $\epsilon > 0$; i.e., $\xi^+ < 1$. Similar calculations for $k < 0$ and $1 < tr^- < e^{-k}$ yield,

$$\begin{aligned}
\int_{x<k} (e^x - 1) \nu(x) dx &= (e^k - 1) \bar{\nu}_{\psi}^-(e^{-k}) - \int_{\mathbb{R}} (e^k - e^x)^+ \nu(x) dx, \\
\int_{\mathbb{R}} (e^k - e^x)^+ \nu(x) dx &= \int_{(0, \infty)} \left(e^k - \frac{1}{x} \right)^+ \nu_{\psi}^-(x) dx = \int_{e^{-k}}^{\infty} \frac{\bar{\nu}_{\psi}^-(x)}{x^2} dx \\
&= \bar{\nu}_{\psi}^-(tr^-) \left(\frac{\xi^-}{\xi^- + 1} \right) (e^k)^{1+1/\xi^-} \left(\frac{\xi^-}{\sigma^-} \right)^{-1/\xi^-} \\
&\quad \times {}_2F_1 \left(1 + \frac{1}{\xi^-}; \frac{1}{\xi^-}; 2 + \frac{1}{\xi^-}; -\frac{1 - tr^- \frac{\xi^-}{\sigma^-}}{e^{-k} \frac{\xi^-}{\sigma^-}} \right), \\
\int_{x<k} (e^x - 1 - x) \nu(x) dx &= \int_{x>e^{-k}} \left(\frac{1}{x} - 1 + \ln x \right) \nu_{\psi}^-(x) dx \\
&= \bar{\nu}_{\psi}^-(e^{-k})(e^k - 1 - k) + \bar{\nu}_{\psi}^-(tr^-) \left(e^k \frac{\sigma^-}{\xi^-} \right)^{1/\xi^-} \xi^- \times K_2, \\
K_2 &= {}_2F_1 \left(\frac{1}{\xi^-}, \frac{1}{\xi^-}, 1 + \frac{1}{\xi^-}; \frac{tr^- \frac{\xi^-}{\sigma^-} - 1}{e^{-k} \frac{\xi^-}{\sigma^-}} \right) - \frac{e^k}{\xi^- + 1} {}_2F_1 \left(1 + \frac{1}{\xi^-}; \frac{1}{\xi^-}; 2 + \frac{1}{\xi^-}; \frac{tr^- \frac{\xi^-}{\sigma^-} - 1}{e^{-k} \frac{\xi^-}{\sigma^-}} \right),
\end{aligned}$$

$$\int_{x < k} x^2 \nu(x) dx = k^2 \bar{\nu}_\psi^-(e^{-k}) + 2 \int_{x > e^{-k}} \frac{\ln(x)}{x} \bar{\nu}_\psi^-(x) dx = k^2 \bar{\nu}_\psi^-(e^{-k}) + 2 \bar{\nu}_\psi^-(tr^-) \times K_3,$$

$$K_3 = e^{k/\xi^-} \xi^- \left(\frac{\xi^-}{\sigma^-} \right)^{-1/\xi^-} \left[\xi^- {}_3F_2 \left(\frac{1}{\xi^-}, \frac{1}{\xi^-}, \frac{1}{\xi^-}; 1 + \frac{1}{\xi^-}, 1 + \frac{1}{\xi^-}; \frac{\xi^- tr^- - 1}{e^{-k} \frac{\xi^-}{\sigma^-}} \right) - k {}_2F_1 \left(\frac{1}{\xi^-}, \frac{1}{\xi^-}, 1 + \frac{1}{\xi^-}; \frac{\xi^- tr^- - 1}{e^{-k} \frac{\xi^-}{\sigma^-}} \right) \right].$$

A.1 Estimation of \mathbb{P} Jump Tails

The estimation builds on the general theoretical results in Bollerslev and Todorov (2010). In particular, let

$$\theta \equiv \left(\sigma^\pm, \xi^\pm, \alpha_0^\pm \bar{\nu}_\psi^\pm(tr^\pm), \alpha_1^\pm \bar{\nu}_\psi^\pm(tr^\pm), \pi \right), \quad \pi \equiv \mathbb{E}(\pi_t),$$

denote the parameter vector that we seek to estimate. Our estimation is then based on the scores associated with the log-likelihood function of the generalized Pareto distribution,

$$\phi_1^\pm(u) = -\frac{1}{\sigma^\pm} + \frac{\xi^\pm u}{(\sigma^\pm)^2} \left(1 + \frac{1}{\xi^\pm} \right) \left(1 + \frac{\xi^\pm u}{\sigma^\pm} \right)^{-1},$$

$$\phi_2^\pm(u) = \frac{1}{(\xi^\pm)^2} \ln \left(1 + \frac{\xi^\pm u}{\sigma^\pm} \right) - \frac{u}{\sigma^\pm} \left(1 + \frac{1}{\xi^\pm} \right) \left(1 + \frac{\xi^\pm u}{\sigma^\pm} \right)^{-1}.$$

The specific estimation equations that we use are,

$$\frac{1}{N} \sum_{t=1}^N \sum_{j=1}^{n-1} \phi_i^\pm \left(\psi^\pm(\Delta_j^{n,t} p) - tr^\pm \right) 1_{\{\psi^\pm(\Delta_j^{n,t} p) > tr^\pm\}} = 0, \quad i = 1, 2,$$

$$\frac{1}{N} \sum_{t=1}^N \sum_{j=1}^{n-1} 1_{\{\psi^\pm(\Delta_j^{n,t} p) > tr^\pm\}} - (1 - \pi) \alpha_0^\pm \bar{\nu}_\psi^\pm(tr^\pm) - \alpha_1^\pm \bar{\nu}_\psi^\pm(tr^\pm) CV_t = 0,$$

$$\frac{1}{N} \sum_{t=2}^N \left(\sum_{j=1}^{n-1} 1_{\{\psi^\pm(\Delta_j^{n,t} p) > tr^\pm\}} - (1 - \pi) \alpha_0^\pm \bar{\nu}_\psi^\pm(tr^\pm) - \alpha_1^\pm \bar{\nu}_\psi^\pm(tr^\pm) CV_t \right) CV_{t-1} = 0,$$

$$(1 - \pi) \frac{1}{N} \sum_{t=1}^N (p_{t+\pi t} - p_t)^2 - \pi \frac{1}{N} \sum_{t=1}^N RV_t = 0.$$

The only remaining issue relates to the choice of tr^\pm . On the one hand, high thresholds are preferred for the asymptotic extreme value approximation in (A.2) to work the best. On the other hand, we also need sufficiently many jumps above the thresholds to make the estimation reliable. Faced with this tradeoff, we set the values for tr^\pm so that they correspond to jumps in the log-prices of 0.6%. In the sample we have 229 such jumps; 127 of

Table 6: Estimates for \mathbb{P} Tail Parameters

The estimates are based on 5-minute high-frequency S&P 500 futures data from January 1990 through June 2007. The truncation levels are fixed at $tr^- = e^{0.006} \approx 1.006$ and $tr^+ = e^{0.006} - 1 \approx 0.006$ for the left and right tails, respectively.

| Left Tail | | | Right Tail | | |
|---|----------|----------|---|----------|----------|
| Parameter | Estimate | St.error | Parameter | Estimate | St.error |
| $\xi_{\mathbb{P}}^-$ | 0.3865 | 0.1050 | $\xi_{\mathbb{P}}^+$ | 0.4361 | 0.1558 |
| $100\sigma_{\mathbb{P}}^-$ | 0.1709 | 0.0222 | $100\sigma_{\mathbb{P}}^+$ | 0.1740 | 0.0272 |
| $\alpha_0^- \bar{\nu}_{\psi}^{\mathbb{P}-}(tr^-)$ | 0.0084 | 0.0030 | $\alpha_0^+ \bar{\nu}_{\psi}^{\mathbb{P}+}(tr^+)$ | 0.0002 | 0.0021 |
| $\alpha_1^- \bar{\nu}_{\psi}^{\mathbb{P}-}(tr^-)$ | 0.0333 | 0.0050 | $\alpha_1^+ \bar{\nu}_{\psi}^{\mathbb{P}+}(tr^+)$ | 0.0361 | 0.0041 |

them negative and 102 of them positive. This choice for tr^{\pm} exceeds the maximum threshold value for $\alpha_{t,i}$ used in the construction of CV_t and JV_t (see equation (B.10) below), so that the continuous component in the price will have negligible effect in our estimation.

The actual results from the estimation are reported in Table 6. As discussed in the main text, the \mathbb{P} tails are clearly time-varying, but otherwise appear close to symmetric. Not reported in the table, our estimate for $\pi \equiv \mathbb{E}(\pi_t)$ equal to 0.2076(0.0448) implies that on average roughly twenty percent of the total variation is due to changes in the price during non-trading hours.

A.2 Estimation of \mathbb{Q} Jump Tails

We start our discussion of the \mathbb{Q} jump tail estimation with a formal proof of Proposition 1.

Proof of Proposition 1: Using the Meyer-Ito formula (Theorem 68 in Protter, 2004) for the payoff function of out-of-the-money calls, the definition of local time (and the behavior of the transitional distribution of a Levy process over small time interval), together with the fact that futures prices are martingales under the risk-neutral measure, we can write for arbitrary $K > F_{t-}$ and $T \downarrow t$,

$$e^{r(t,T)} C_t(K) \sim \int_t^T \mathbb{E}_t^{\mathbb{Q}} \left(\int_{\mathbb{R}} [1_{\{F_s < K\}} (F_s - e^x - K)^+ + 1_{\{F_s > K\}} (K - F_s - e^x)^+] \nu_s^{\mathbb{Q}}(dx) \right) ds, \quad (\text{A.4})$$

where $f(x) \sim g(x)$ is to be interpreted as $\lim_{x \rightarrow 0} \frac{f(x)}{g(x)} = 1$. A similar result holds for close-to-maturity out-of-the-money puts. This result was first derived in Carr and Wu (2003) in the proof of their Theorem 1.

Next, note that since F_t is a martingale under \mathbb{Q} , using Doob's inequality for $K > F_{t-}$, i.e., for call options out-of-the-money, it follows that

$$\mathbb{Q}_t \left(\sup_{u \in [t, T]} |F_u - F_{t-}| \geq K - F_{t-} \right) \leq C \times (T - t), \quad (\text{A.5})$$

where the constant C is adapted to the information set at time t , and is otherwise proportional to the risk-neutral, conditional to time t , expectation of the quadratic variation of the \mathbb{Q} -martingale process $\{F_t\}_t$. With this, we may therefore write

$$\begin{aligned} e^{r(t,T)}C_t(K) &= \int_t^T \mathbb{E}_t^{\mathbb{Q}} \int_{\mathbb{R}} 1_{\{F_{s-} < K\}} (F_{s-}e^x - K)^+ \nu_s^{\mathbb{Q}}(dx) ds + O((T-t)^2) \\ &= \int_t^T \mathbb{E}_t^{\mathbb{Q}} 1_{\{F_{t-} < K\}} \int_{\mathbb{R}} (F_{s-}e^x - K)^+ \nu_s^{\mathbb{Q}}(dx) ds + O((T-t)^2). \end{aligned} \quad (\text{A.6})$$

To continue further, note that for every $x \in \mathbb{R}$,

$$|(F_{s-}e^x - K)^+ - (F_{t-}e^x - K)^+| \leq e^x |F_{s-} - F_{t-}|. \quad (\text{A.7})$$

Now, making use of the Burkholder-Davis-Gundy inequality it follows that

$$\mathbb{E}_t^{\mathbb{Q}} \left(\sup_{u \in [t, T]} |F_u - F_{t-}| \right) \leq C\sqrt{T-t}, \quad (\text{A.8})$$

where the constant C satisfies the same restrictions spelled out immediately above. Combining these results, we have

$$e^{r(t,T)}C_t(K) = (T-t)F_{t-}RT_t^{\mathbb{Q}}(k) + O((T-t)^{3/2}), \quad (\text{A.9})$$

where $RT_t^{\mathbb{Q}}(k)$ is formally defined in equation (2.4).

Using the decomposition of the risk-neutral Lévy density in (2.6), it follows further that

$$\frac{RT_t^{\mathbb{Q}}(k_1)}{RT_t^{\mathbb{Q}}(k_2)} = \frac{\int_{\mathbb{R}} (e^x - e^{k_1})^+ \nu_{\mathbb{Q}}(x) dx}{\int_{\mathbb{R}} (e^x - e^{k_2})^+ \nu_{\mathbb{Q}}(x) dx},$$

where $k_1 = \frac{K_1}{F_{t-}}$ and $k_2 = \frac{K_2}{F_{t-}}$. Finally, applying the above formulas with $tr^+ = K_1/F_{t-} - 1$, together with the result in Smith (1987) concerning the residual function $\tau^+(x)$, concludes the proof for the right tail. The proof for the puts and the left tail proceeds in a completely analogous fashion. \square

In the proof immediately above we used a somewhat stronger version of the approximation in (A.2), and the fact that the true value of σ^{\pm} is ξ^{\pm}/x . A similar restriction underlies the popular Hill estimator of the tail index. However, this stronger version of (A.2) is not scale invariant. Therefore, while Proposition 1 conveys the main idea, in our actual estimation we only make use of (A.2).

In particular, recall that our estimation of the jump tail probabilities under \mathbb{Q} only requires the weak separability condition for the time dependence and jump sizes implicit in the Lévy density in equation (2.6). The parameter vector that we seek to estimate is

$$\theta \equiv \left(\sigma_{\mathbb{Q}}^{\pm}, \xi_{\mathbb{Q}}^{\pm}, \alpha_{\mathbb{Q}}^{\pm} \bar{\nu}_{\psi}^{\mathbb{Q}^{\pm}}(tr^{\pm}) \right).$$

From the calculations above, we have that for $0 < tr^+ \leq e^k - 1$,

$$\mathbb{E} \left(RT_t^{\mathbb{Q}}(k) \right) = \alpha_{\mathbb{Q}}^+ \bar{\nu}_{\psi}^{\mathbb{Q}^+}(tr^+) \frac{\sigma_{\mathbb{Q}}^+}{1 - \xi_{\mathbb{Q}}^+} \left(1 + \frac{\xi_{\mathbb{Q}}^+}{\sigma_{\mathbb{Q}}^+} (e^k - 1 - tr^+) \right)^{1 - 1/\xi_{\mathbb{Q}}^+},$$

while for $1 < tr^- \leq e^{-k}$,

$$\begin{aligned} \mathbb{E} \left(LT_t^{\mathbb{Q}}(k) \right) &= \alpha_{\mathbb{Q}}^- \bar{\nu}_{\psi}^{\mathbb{Q}^-}(tr^-) \frac{\xi_{\mathbb{Q}}^-}{\xi_{\mathbb{Q}}^- + 1} (e^k)^{1 + 1/\xi_{\mathbb{Q}}^-} \left(\frac{\xi_{\mathbb{Q}}^-}{\sigma_{\mathbb{Q}}^-} \right)^{-1/\xi_{\mathbb{Q}}^-} \\ &\quad \times {}_2F_1 \left(1 + \frac{1}{\xi_{\mathbb{Q}}^-}; \frac{1}{\xi_{\mathbb{Q}}^-}; 2 + \frac{1}{\xi_{\mathbb{Q}}^-}; \frac{tr^- \frac{\xi_{\mathbb{Q}}^-}{\sigma_{\mathbb{Q}}^-} - 1}{e^{-k} \frac{\xi_{\mathbb{Q}}^-}{\sigma_{\mathbb{Q}}^-}} \right), \end{aligned}$$

where $\alpha_{\mathbb{Q}}^+ \equiv \mathbb{E} \left(\frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \varphi_s^+ ds \right) \right)$ and $\alpha_{\mathbb{Q}}^- \equiv \mathbb{E} \left(\frac{1}{T-t} \mathbb{E}_t^{\mathbb{Q}} \left(\int_t^T \varphi_s^- ds \right) \right)$.

Now replacing the unconditional risk-neutral expectations of the jump tail measures with the sample analogues of $RT_t^{\mathbb{Q}}(k)$ and $LT_t^{\mathbb{Q}}(k)$ for three different levels of moneyness for each of the tails results in an exactly identified GMM estimator. The levels of moneyness that we use in the estimation are $\frac{K}{F_{t-}} = 0.9250, 0.9125$, and 0.9000 for the left tail, and $\frac{K}{F_{t-}} = 1.0750, 1.0875$ and 1.1000 for the right tail. These levels of moneyness are sufficiently “deep” in the tails to guarantee that the effect of the diffusive price component in measuring $LT_t^{\mathbb{Q}}(k)$ and $RT_t^{\mathbb{Q}}(k)$ is minimal, and that the extreme value distribution provides a good approximation to the jump tail probabilities. Also, the truncation levels are fixed at $tr^- = (1 - 0.075)^{-1}$ and $tr^+ = 0.075$ for the left and right tails, respectively. This corresponds directly to the moneyness for the closest-to-the-money options used in the estimation.

Table 7: Estimates for \mathbb{Q} Tail Parameters

The estimates are based on S&P 500 options data from January 1996 through June 2007. The log-moneyness of the options used for the left and right tails are fixed at $k = 0.9$ and $k = 1.1$, respectively.

| Left Tail | | | Right Tail | | |
|--|----------|----------|--|----------|----------|
| Parameter | Estimate | St.error | Parameter | Estimate | St.error |
| $\xi_{\mathbb{Q}}^-$ | 0.2581 | 0.02816 | $\xi_{\mathbb{Q}}^+$ | 0.0793 | 0.0147 |
| $\sigma_{\mathbb{Q}}^-$ | 0.0497 | 0.0021 | $\sigma_{\mathbb{Q}}^+$ | 0.0238 | 0.0010 |
| $\alpha_{\mathbb{Q}}^- \bar{\nu}_{\psi}^{\mathbb{Q}^-} \left(\frac{1}{1-0.075} \right)$ | 0.9888 | 0.0525 | $\alpha_{\mathbb{Q}}^+ \bar{\nu}_{\psi}^{\mathbb{Q}^+}(0.075)$ | 0.5551 | 0.0443 |

The results from the estimation are reported in Table 7. Not only is the left tail under the \mathbb{Q} distribution at a higher level, it also decays at a slower rate. Although not directly comparable, the standard errors for the \mathbb{Q} tail parameters are much smaller than the ones for the \mathbb{P} tail parameters reported in Table 6 above. This is quite intuitive as our estimation for the risk-neutral distribution are based on expectations for the “large” jumps, while the

estimates for the statistical distribution rely on actual jump realizations in the high-frequency data for inferring the relevant expectations, thus involving an extra layer of estimation error uncertainty.

B Estimation of Realized Measures

This appendix provides further details concerning the estimation of CV_t , RJV_t and LJV_t in equations (3.2)-(3.5) and the relevant truncation levels. Theoretically, any $\alpha > 0$ and $\varpi \in (0, 0.5)$ will work. In the actual estimation, however, we fix $\varpi = 0.49$. Our choice of α is a bit more involved and takes into account that volatility changes across days and also has a strong diurnal pattern over the trading day.

To account for the latter feature we estimate nonparametrically a time-of-day factor TOD_i , $i=1,2,\dots,n$,

$$TOD_i = NOI_i \frac{\sum_{t=1}^N (f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}})^2 \mathbf{1}_{\{|f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}}| \leq \bar{\alpha}\Delta_n^{0.49}\}}}{\sum_{t=1}^N \sum_{i=1}^M (f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}})^2}, \quad (\text{B.10})$$

where

$$NOI_i = \frac{\sum_{t=1}^N \sum_{i=1}^{n-1} \mathbf{1}_{\{|f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}}| \leq \bar{\alpha}\Delta_n^{0.49}\}}}{\sum_{t=1}^N \mathbf{1}_{\{|f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}}| \leq \bar{\alpha}\Delta_n^{0.49}\}}},$$

$$\bar{\alpha} = 3\sqrt{\frac{\pi}{2}} \sqrt{\frac{1}{N} \sum_{t=1}^N \sum_{i=2}^{n-1} |f_{(t-1+\pi_t)+i\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}}| |f_{(t-1+\pi_t)+(i-1)\Delta_{n,t}} - f_{(t-1+\pi_t)+(i-2)\Delta_{n,t}}|},$$

and $\Delta_n \equiv \frac{1}{n}$. The truncation level $\bar{\alpha}$ is based on the average volatility in the sample, as measured by the so-called bipower variation measure defined in Barndorff-Nielsen and Shephard (2004). Consequently $\bar{\alpha} \approx \sqrt{(1 - \mathbb{E}(\pi_t))\mathbb{E}(\sigma_t^2)}$. The NOI_i factor puts the numerator and denominator in equation (B.10) in the same units. This factor is approximately equal to $n - 1$. This calculation of TOD_i implicitly assumes that π_t and σ_t^2 are both stationary and ergodic processes. The resulting TOD_i estimates are plotted in Figure 6. The figure shows the previously well-documented U-shape in the average volatility over NYSE trading hours, with separate end effects for the futures before and after the cash market opens; see, e.g., the related discussion in Andersen and Bollerslev (1997) pertaining to an earlier sample period.

To account for the time-varying volatility across days, we use the estimated continuous volatility for the previous day (for the first day in the sample we use $\bar{\alpha}$). Putting all of the pieces together, our time-varying threshold may be succinctly expressed as $\alpha_{t,i} = 3\sqrt{CV_{t-1}} \times TOD_i \times \Delta_n^{0.49}$.

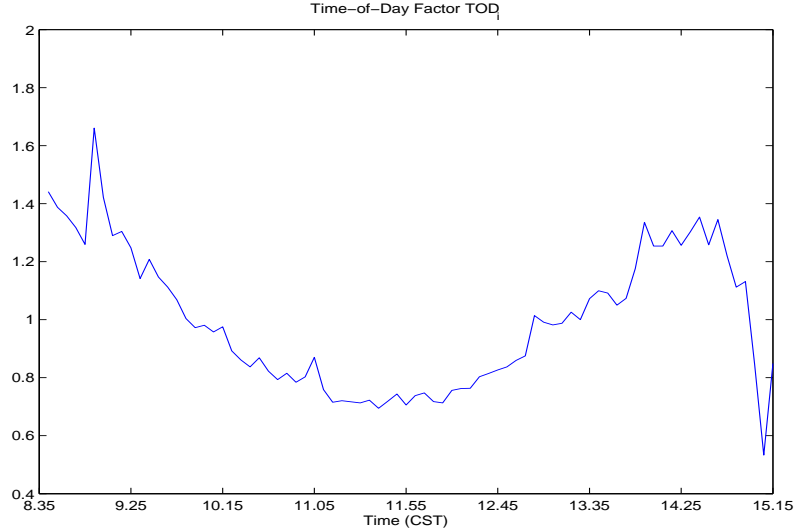


Figure 6: *Time-of-Day Factor*. The estimates are based on 5-minute high-frequency S&P 500 futures data from January 1990 through June 2007.

C Estimation of Expected Integrated Volatility

Our estimation of $\mathbb{E}_t \left(\int_t^T \sigma_s^2 ds \right)$ is based on the following restricted 22-order Vector Autoregression,

$$\mathbf{X}_t = \mathbf{A}_0 + \mathbf{A}_1 \mathbf{X}_{t-1} + \mathbf{A}_5 \sum_{i=1}^5 \mathbf{X}_{t-i}/5 + \mathbf{A}_{22} \sum_{i=1}^{22} \mathbf{X}_{t-i}/22 + \epsilon_t, \quad (\text{C.11})$$

for the four-dimensional vector

$$\mathbf{X}_t \equiv (CV_t \quad RJV_t \quad LJV_t \quad (p_{t+\pi_t} - p_t)^2)'$$

We approximate the multi-period expected integrated volatility from the sum of the relevant forecasts for the first element in \mathbf{X}_t , which we in turn approximate by easy-to-calculate recursive linear projections using the Kalman filter. Similar reduced-form forecasting procedure for realized variation measures have been demonstrated to work well empirically by Andersen et al. (2003) among many others.

The resulting parameter estimates for the VAR are reported in Table 8. Not surprisingly, the estimates for the (1, 1) elements in the \mathbf{A}_1 , \mathbf{A}_5 and \mathbf{A}_{22} matrices point to the existence of strong own persistence in the continuous volatility component. The estimates also reveal significant asymmetries in the dynamic dependencies between the continuous volatility and the tails and the overnight increments, underscoring the importance of including these variables in the construction of the integrated volatility forecasts.

Table 8: Estimation Results for the restricted VAR(22) model in (C.11)

The realized variation measures underlying the estimates are based on 5-minute high-frequency S&P 500 futures data from January 1990 through June 2007. Standard errors are reported in parentheses.

| | | | | |
|-------------------|------------------|------------------|------------------|------------------|
| \mathbf{A}'_0 | 0.0368 (0.0190) | 0.0010 (0.0216) | 0.0430 (0.0188) | -0.0153 (0.0456) |
| \mathbf{A}_1 | 0.6956 (0.0316) | -0.0728 (0.0324) | 0.2022 (0.0380) | -0.0618 (0.0146) |
| | -0.0449 (0.0360) | -0.0348 (0.0369) | 0.1145 (0.0432) | -0.0299 (0.0167) |
| | -0.0028 (0.0313) | -0.0175 (0.0321) | 0.0367 (0.0376) | -0.0211 (0.0145) |
| | 0.0831 (0.0758) | -0.0250 (0.0778) | 0.1788 (0.0910) | -0.0613 (0.0351) |
| \mathbf{A}_5 | 0.0672 (0.0462) | 0.0325 (0.0841) | -0.0974 (0.0947) | 0.0709 (0.0341) |
| | 0.0734 (0.0525) | -0.0515 (0.0957) | 0.0018 (0.1077) | 0.0532 (0.0388) |
| | -0.0001 (0.0457) | 0.0018 (0.0833) | 0.0063 (0.0938) | 0.0512 (0.0338) |
| | 0.2325 (0.1107) | 0.0495 (0.2016) | -0.1477 (0.2270) | 0.2938 (0.0818) |
| \mathbf{A}_{22} | 0.1669 (0.0464) | -0.2100 (0.1786) | 0.0962 (0.1857) | 0.0285 (0.0606) |
| | 0.0399 (0.0528) | -0.2202 (0.2031) | 0.3187 (0.2113) | 0.0353 (0.0689) |
| | 0.0308 (0.0460) | -0.0964 (0.1769) | 0.1724 (0.1840) | 0.0386 (0.0600) |
| | -0.0097 (0.1113) | -0.0355 (0.4280) | 0.6235 (0.4452) | -0.1462 (0.1452) |

D Rank Test

We briefly sketch the basic ideas and calculations required for the rank tests that form the basis for assessing Propositions 2 and 3; additional details concerning the implementation of the test are given in Gill and Lewbel (1992). Let A denote a general $p \times p$ matrix. The hypothesis of interest may then be expressed as,

$$H_0 : \text{rank}(A) \equiv r \quad \text{vs.} \quad H_1 : \text{rank}(A) > r, \quad r < p,$$

where we have an asymptotically normally distributed estimator of A , say

$$\sqrt{N} \text{vec}(\hat{A} - A) \xrightarrow{\mathbb{L}} \mathcal{N}(0, V).$$

Also, let \hat{V} denote a consistent estimator of V . The rank test is based on LDU (lower-diagonal-upper triangular) decomposition with complete pivoting; i.e., the elements of the diagonal matrix are decreasing in absolute value. To this end, write the matrices associated with the LDU decomposition as

$$\hat{P}\hat{A}\hat{Q} = \hat{L}\hat{D}\hat{U},$$

where \hat{P} and \hat{Q} are permutation matrixes,

$$\hat{L} = \begin{pmatrix} \hat{L}_{11} & \mathbf{0} \\ \hat{L}_{21} & \mathbf{I}_{p-r} \end{pmatrix} \quad \hat{D} = \begin{pmatrix} \hat{D}_1 & \mathbf{0} \\ \mathbf{0} & \hat{D}_2 \end{pmatrix} \quad \hat{U} = \begin{pmatrix} \hat{U}_{11} & \hat{U}_{12} \\ \mathbf{0} & \hat{U}_{22} \end{pmatrix}$$

and \widehat{L}_{11} , \widehat{U}_{11} and \widehat{U}_{22} are unit lower triangular. In addition define,

$$\widehat{H} = [-\widehat{L}_{22}^{-1}\widehat{L}_{21}\widehat{L}_{11}^{-1} \quad \widehat{L}_{22}^{-1} \quad \mathbf{0}] \quad \widehat{K} = \begin{pmatrix} -\widehat{U}_{11}^{-1}\widehat{U}_{12}\widehat{U}_{22}^{-1} \\ \widehat{U}_{22}^{-1} \end{pmatrix}.$$

Further, let Ξ_i denote the $(p-r) \times (p-r)$ matrix with ones along the i -th diagonal and zeros everywhere else, and define

$$\Xi = \begin{pmatrix} \Xi_1 \\ \vdots \\ \Xi_n \end{pmatrix},$$

and

$$\widehat{W} = \Xi'(\widehat{K}' \otimes \widehat{H})(\widehat{Q}' \otimes \widehat{P})\widehat{V}(\widehat{Q} \otimes \widehat{P}')(\widehat{K} \otimes \widehat{H}')\Xi.$$

Finally, let $\widehat{d}_2 = \text{diag}(\widehat{D}_2)$. The rank test statistic,

$$\chi = N\widehat{d}_2'\widehat{W}^{-1}\widehat{d}_2,$$

is then distributed as χ_{p-r}^2 under the null hypothesis that $\text{rank}(A) \equiv r$.

References

- Aït-Sahalia, Y. and A. Lo (1998). Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. *Journal of Finance* 53, 499–548.
- Aït-Sahalia, Y. and A. Lo (2000). Nonparametric Risk Management and Implied Risk Aversion. *Journal of Econometrics* 94, 9–51.
- Aït-Sahalia, Y., Y. Wang, and F. Yared (2001). Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset? *Journal of Econometrics* 102, 67–110.
- Andersen, T. G., L. Benzoni, and J. Lund (2002). An Empirical Investigation of Continuous-Time Equity Return Models. *Journal of Finance* 57, 1239–1284.
- Andersen, T. G. and T. Bollerslev (1997). Intraday Periodicity and Volatility Persistence in Financial Markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T. G. and T. Bollerslev (1998). Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2007). Roughing it up: Disentangling Continuous and Jump Components in Measuring, Modeling and Forecasting Asset Return Volatility. *Review of Economics and Statistics* 73, 251–277.

- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and Forecasting Realized Volatility. *Econometrica* 71, 579–625.
- Andersen, T. G., T. Bollerslev, and N. Meddahi (2004). Analytic Evaluation of Volatility Forecasts. *International Economic Review* 45, 1079–1110.
- Andersen, T. G. and O. Bondarenko (2010). Dissecting the Market Pricing of Volatility. Working paper, Northwestern University.
- Anderson, E., L. P. Hansen, and T. Sargent (2003). A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1, 68–123.
- Artzner, P., F. Delbaen, J. M. Eber, and D. Heath (1999). Coherent Measures of Risk. *Mathematical Finance* 9, 203–228.
- Bakshi, G., C. Cao, and Z. Chen (1997). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52, 2003–2049.
- Bakshi, G. and N. Kapadia (2003). Delta-Hedged Gains and the Negative Market Volatility Risk Premium. *Review of Financial Studies* 16, 527–566.
- Bakshi, G., N. Kapadia, and D. Madan (2003). Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies* 16, 101–143.
- Bakshi, G. and D. Madan (2000). Spanning and Derivative Security Valuation. *Journal of Financial Economics* 55, 205–238.
- Bakshi, G. and D. Madan (2006). A Theory of Volatility Spreads. *Management Science* 52, 1945–1956.
- Bakshi, G., D. Madan, and G. Panayotov (2009). Deducing the Implications of Jump Models for the Structure of Stock Market Crashes, Rallies, Jump Arrival Rates, and Extremes. *Journal of Business and Economic Statistics* 28, 380–396.
- Bansal, R. and I. Shaliastovich (2009). Confidence Risk and Asset Prices. Working paper, Duke University.
- Bansal, R. and A. Yaron (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59, 1481–1509.
- Barndorff-Nielsen, O. E., S. Kinnebrock, and N. Shephard (2010). Measuring Downside Risk - Realised Semivariance. In T. Bollerslev, J. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics, Essays in Honor of Robert F. Engle*. Oxford University Press.
- Barndorff-Nielsen, O. E. and N. Shephard (2001). Non- Gaussian Ornstein-Uhlenbeck-based Models and Some of Their Applications in Financial Economics. *Journal of the Royal Statistical Society: Series B* 63, 167–241.

- Barndorff-Nielsen, O. E. and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barro, R. J. (2006). Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics* 121, 823–866.
- Bates, D. S. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies* 9, 69–107.
- Bates, D. S. (2000). Post-'87 Crash Fears in S&P 500 Future Options. *Journal of Econometrics* 94, 181–238.
- Bates, D. S. (2008). The Market for Crash Risk. *Journal of Economic Dynamics and Control* 32, 2291–2321.
- Bates, D. S. (2009). U. S. Stock Market Crash Risk, 1926–2006. Working paper, University of Iowa.
- Bollerslev, T., U. Kretschmer, C. Pigorsch, and G. Tauchen (2009). A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects. *Journal of Econometrics* 150, 151–166.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T. and V. Todorov (2010). Estimation of Jump Tails. Working paper, Duke and Northwestern Universities.
- Bondarenko, O. (2003). Why are Put Options so Expensive? Working paper, University of Illinois, Chicago.
- Britten-Jones, M. and A. Neuberger (2000). Option Prices, Implied Price Processes, and Stochastic Volatility. *Journal of Finance* 55, 839–866.
- Broadie, M., M. Chernov, and M. Johannes (2009a). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance* 62, 1453–1490.
- Broadie, M., M. Chernov, and M. Johannes (2009b). Understanding Index Options Returns. *Review of Financial Studies* 22, 4493–4529.
- Campbell, J. and J. Cochrane (1999). By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107, 205–251.
- Carr, P., H. Geman, D. B. Madan, and M. Yor (2003). Stochastic Volatility for Lévy Processes. *Journal of Business* 75, 345–382.
- Carr, P. and L. Wu (2003). What Type of Process Underlies Options? A Simple Robust Test. *Journal of Finance* 58, 2581–2610.
- Carr, P. and L. Wu (2009). Variance Risk Premia. *Review of Financial Studies* 22, 1311–1341.

- Chacko, G. and L. Viceira (2005). Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets. *Review of Financial Studies* 18, 1369–1402.
- Cochrane, J. (2005). *Asset Pricing* (2nd ed.). Princeton University Press.
- Cont, R. and P. Tankov (2009). Constant Proportion Portfolio Insurance in the Presence of Jumps in Asset Prices. *Mathematical Finance* 19, 379–401.
- Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics* 7, 174–196.
- Delbaen, F. and W. Schachermayer (1994). A General Version of the Fundamental Theorem of Asset Pricing. *Mathematische Annalen* 300, 520–563.
- Drechsler, I. (2011). Uncertainty, Time-Varying Fear, and Asset Prices. Working paper, New York University.
- Drechsler, I. and A. Yaron (2011). What’s Vol Got to Do with It? *Review of Financial Studies* 24, 1–45.
- Duffie, D. (2001). *Dynamic Asset Pricing Theory* (3rd ed.). Princeton University Press.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.
- Embrechts, P., C. Kluppelberg, and T. Mikosch (2001). *Modelling Extremal Events* (3rd ed.). Berlin: Springer-Verlag.
- Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Eraker, B. (2008). The Volatility Premium. Working paper, University of Wisconsin.
- Foresi, S. and L. Wu (2005). Crash-o-phobia: A Domestic Fear or a Worldwide Phenomenon? *Journal of Derivatives* 9, 8–21.
- Gabaix, X. (2010). Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance. Working paper, New York University.
- Gill, L. and A. Lewbel (1992). Testing the Rank and Definiteness of Estimated Matrices with Applications to Factor, State-Space and ARMA Models. *Journal of the American Statistical Association* 87, 766–776.
- Jacod, J. and A. N. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.
- Jiang, G. and Y. Tian (2005). The Model-Free Implied Volatility and Its Information Content. *Review of Financial Studies* 18, 1305–1342.
- Jiang, G. and Y. Tian (2007). Extracting Model-Free Volatility from Option Prices: An Estimation the VIX Index. *Journal of Derivatives* 14, 1–26.

- Mancini, C. (2001). Disentangling the Jumps of the Diffusion in a Geometric Brownian Motion. *Giornale dell'Istituto Italiano degli Attuari LXIV*, 19–47.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics 63*, 3–50.
- Pan, J. and J. Liu (2003). Dynamic Derivative Strategies. *Journal of Financial Economics 69*, 401–430.
- Pan, J., J. Liu, and F. Longstaff (2003). Dynamic Asset Allocation with Event Risk. *Journal of Finance 58*, 231–259.
- Pan, J., J. Liu, and T. Wang (2005). An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks. *Review of Financial Studies 18*, 131–164.
- Pozdnyakov, V. and J. Steele (2004). On the Martingale Framework for Futures Prices. *Stochastic Processes and their Applications 109*, 69–77.
- Protter, P. (2004). *Stochastic Integration and Differential Equations* (2nd ed.). Berlin: Springer-Verlag.
- Rietz, T. (1988). The Equity Risk Premium: A Solution. *Journal of Monetary Economics 22*, 117–131.
- Rosenberg, J. and R. F. Engle (2002). Empirical Pricing Kernels. *Journal of Financial Economics 64*, 341–372.
- Rubinstein, M. (1994). Implied Binomial Trees. *Journal of Finance 69*, 771–817.
- Santa-Clara, P. and S. Yan (2010). Crashes, Volatility and the Equity Premium: Lessons from S&P 500 Options. *Review of Economics and Statistics 92*, 435–451.
- Singleton, K. (2006). *Empirical Dynamic Asset Pricing*. Princeton University Press.
- Sizova, N. (2009). Integrated Variance Forecasting: Model-Based vs. Reduced-Form. Working paper, Rice University.
- Smith, R. (1987). Estimating Tails of Probability Distributions. *Annals of Statistics 15*, 1174–1207.
- Tankov, P. (2009). Pricing and Hedging Gap Risk. Working paper, Ecole Polytechnique.
- Todorov, V. (2010). Variance Risk Premia Dynamics: The Role of Jumps. *Review of Financial Studies 23*, 345–383.
- Todorov, V. and G. Tauchen (2011). Volatility Jumps. *Journal of Business and Economic Statistics*, forthcoming.
- Wachter, J. A. (2010). Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? Working paper, University of Pennsylvania.
- Whaley, R. (2009). Understanding VIX. *Journal of Portfolio Management 35*, 98–105.

Zhang, L., P. Mykland, and Y. Aït-Sahalia (2005). A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data. *Journal of the American Statistical Association* 100, 1394–1411.