The Risk Premia Embedded in Index Options*

Torben G. Andersen† Nicola Fusari‡ Viktor Todorov§

October 12, 2013

Abstract

We study the dynamic relation between aggregate stock market risks and risk premia via an exploration of the time series of equity-index option surfaces. The analysis is based on estimating a general parametric asset pricing model for the risk-neutral equity market dynamics using a panel of options on the S&P 500 index, while remaining fully nonparametric about the actual evolution of market risks. We find that the risk-neutral jump intensity, which controls the pricing of left tail risk, cannot be spanned by the market volatility (and its components), so an additional factor is required to account for its dynamics. This tail factor has no incremental predictive power for future equity return volatility or jumps beyond what is captured by the current and past level of volatility. In contrast, the novel factor is critical in predicting the future market excess returns over horizons up to one year, and it explains a large fraction of the future variance risk premium. We contrast our findings with those implied by structural asset pricing models that seek to rationalize the predictive power of option data. Relative to those studies, our findings suggest a wider wedge between the dynamics of equity market risks and the corresponding risk premia with the latter typically displaying a far more persistent reaction following market crises.

Keywords: Option Pricing, Risk Premia, Jumps, Stochastic Volatility, Return Predictability, Risk Aversion, Extreme Events.

JEL classification: C51, C52, G12.

*Andersen gratefully acknowledges support from CREATES funded by the Danish National Research Foundation. Todorov's work was partially supported by NSF Grant SES-0957330. We would like to thank Snehal Banerjee, Peter Carr, Anna Cieslak, Bjorn Eraker, Ravi Jaganathan, Robert Merton, Sergio Rebelo, Myron Scholes, Ivan Shaliastovich, Liuren Wu as well as seminar participants at Kellogg School of Management, Northwestern University, Duke University, the "Montreal 2013 Econometrics Conference: Time Series and Financial Econometrics," the 40$^{\text{th}}$ Annual Meeting of the Danish Econometric Society, Sandbjerg, Denmark, the European University Institute, Florence, 2013 Workshop on "Measuring and Modeling Financial Risk and High Frequency Data", the 2013 "International Conference in Financial Econometrics" at Shandong University, Jinan, China and the "40 Years after the Black-Scholes-Merton Model" conference at the Stern School of Business, October 2013, for helpful comments.

†Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: t-andersen@northwestern.edu.

‡The Johns Hopkins University Carey Business School, Baltimore, MD 21202; e-mail: nicola.fusari@jhu.edu.

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

1 Introduction

Equity markets are subject to pronounced time-variation in volatility as well as abrupt shifts, or jumps. Moreover, these risk features are related in intricate ways, inducing a complex equity return dynamics. Hence, the markets are incomplete and derivative securities, written on the equity index, are non-redundant assets. This partially rationalizes the rapid expansion in the trading of contracts offering distinct exposures to volatility and jump risks. From an economic perspective, it suggests that derivatives data contain important information regarding the risk and risk pricing of the underlying asset. Indeed, recent evidence, exploiting parametric models, e.g., Broadie et al. (2009), or nonparametric techniques, e.g., Bollerslev and Todorov (2011b), finds the pricing of jump risk, implied by option data, to account for a significant fraction of the equity risk premium.

Standard no-arbitrage and equilibrium-based asset pricing models imply a tight relationship between the dynamics of the options and the underlying asset. This arises from the assumptions concerning the pricing of risk in the no-arbitrage setting and the endogenous pricing kernels implied by the equilibrium models. A prominent example is the illustrative double-jump model of Duffie et al. (2000) in which the return volatility itself follows an affine jump diffusion. In this context, the entire option surface is governed by the evolution of market volatility, i.e., the dynamics of all options is driven by a single latent Markov (volatility) process.

Recent empirical evidence reveals, however, that the dynamics of the option surface is far more complex. For example, the term structure of the volatility index, VIX, shifts over time in a manner that is incompatible with the surface being driven by a single factor. Likewise, Bates (2000) documents that a two-factor stochastic volatility model for the risk-neutral market dynamics provides a significant improvement over a one-factor version. Moreover, Bollerslev and Todorov (2011b) find that even the short-term option dynamics cannot be captured adequately by a single factor as the risk-neutral tails display independent variation relative to market volatility, thus driving a wedge between the dynamics of the option surface and the underlying asset prices.

The objective of the current paper is to characterize the risk premia, implied by the large panel of S&P 500 index options, and its relation with the aggregate market risks in the economy. As discussed in Andersen et al. (2012), the option panel contains rich information both for the evolution of volatility and jump risks and their pricing. Consequently, we let the option data speak for themselves in determining the risk premium dynamics and discriminating among alternative hypotheses regarding the source of variation in risk as well as risk pricing.

The standard no-arbitrage approach starts by estimating a parametric model characterizing the evolution of the underlying asset price. Risk premia are then introduced through a pricing

kernel which implies that risk compensation is obtained through parameter shifts. This ensures, conveniently, that the associated risk-neutral dynamics remains within the same parametric class entertained for the statistical measure, see, e.g., Singleton (2006), Chapter 15. However, this approach tends to tie the equity market and option surface dynamics closely together. In particular, the equity risk premia are typically linear in (components of) volatility. In contrast, we find the options to display risk price variation that is largely unrelated to, and effectively unidentifiable from, the underlying asset prices alone.

This motivates our "reverse" approach of starting with a parametric model for the risk-neutral dynamics and estimating it exclusively from option data along with no-arbitrage restrictions based on model-free volatility measures constructed from the underlying asset data. In this manner, we avoid letting a (possibly misspecified) parametric structure for the $\mathbb{P}$-dynamics impact the identification of option risk premia. We then explore the risk premia dynamics by combining the extracted state vector with high-frequency data on the equity index.

We document that even a very general two-factor stochastic volatility model with jumps both in price and volatility as well as a time-varying jump intensity produces systematic biases in the fit to the option surface. The problems are particularly acute following periods like the Asian crisis in 1997 or the great recession originating in the Fall of 2008. After these events, market volatility reverts back towards its pre-crisis level fairly quickly, while out-of-the-money put options remain expensive, inducing a steepening of the (Black-Scholes) implied volatility curve. This type of variation in the option surface is difficult to accommodate for standard (no-arbitrage or equilibrium) models as they imply that the priced jump tail risk, which in turn determines the out-of-the-money put prices, is governed solely by (components of) market volatility. This feature implies, in particular, that these models tend to generate risk premia that are too low in the aftermath of crises.

We follow Andersen et al. (2012) by introducing a third factor driving the risk-neutral jump intensity. This novel tail factor is not part of the volatility dynamics although it may be correlated with the level of volatility. We model it as purely jump driven, with one component jointly governed by the volatility jumps while another is independent of the volatility process. This feature allows the jump intensity to escalate – through so-called self-excitation of the jumps – in periods of crises when price and volatility jumps are prevalent, thus "magnifying" the response of the jump intensity to major (negative) market shocks. The extended model remains within the popular class of affine jump-diffusion models of Duffie et al. (2000) and exemplifies the flexibility of such models for generating intricate, yet analytically tractable, dynamic interactions between volatility and jump risks. This type of extension has not been explored in prior empirical option pricing studies.

Relative to the empirical analysis in Andersen et al. (2012), we estimate the model by minimizing, not the squared, but the *relative* squared option pricing error across the full sample. This reduces the weight assigned to highly turbulent periods where the bid-ask spreads and pricing errors increase sharply. We find that the tail factor improves the characterization of the option surface dynamics very significantly. In particular, the new model no longer undervalues short-maturity out-of-the-money puts in the aftermaths of crises. Hence, our extended risk-neutral model provides a more suitable basis for studying the dynamic properties of market risk premia.

In turn, the presence of this independently evolving tail factor implies that part of the risk premium dynamics cannot be captured by the state variables driving the underlying asset price dynamics. This implies that this jump risk factor may have predictive power for future risk premia over and above what is implied by the volatility factors. This is indeed what we find. The novel tail factor is significant in forecasting future excess market returns for horizons up to one year while the volatility factors are insignificant. Similarly, the new tail factor is important for predicting the future variance risk premium in conjunction with one of the volatility components. Taken together, our findings rationalize why the variance risk premium provides superior forecasts for future returns relative to volatility itself, as documented in Bollerslev et al. (2009). The key is the existence of the new factor driving the left jump tails of the risk-neutral distribution.

Importantly, while the new jump factor has predictive power for risk premia, it contains no incremental information regarding the future evolution of volatility and jump risks for the underlying asset relative to the traditional volatility factors. Hence, our findings indicate that option markets embody critical information about the market risk premia and its dynamics which is essentially unidentifiable from stock market data alone. Moreover, the option surface dynamics contains information that can improve the modeling and forecasting of future volatility and jump risks, but such applications necessitate an initial untangling of the components in the risk premia that evolve independently from the volatility process. Overall, our empirical results suggest that there is a wedge between the stochastic evolution of risks in the economy and their pricing, with the latter typically having a far more persistent response to (negative) tail events than the former.

Our finding of a substantial wedge between the dynamics of the option and stock markets presents a challenge for traditional structural asset pricing models. Specifically, the standard exponentially-affine equilibrium models with a representative agent equipped with Epstein-Zin preferences imply that the ratio of the risk-neutral and statistical jump tails is constant. On the contrary, the new factor, extracted from the option data, drives the risk-neutral jump tail but has no discernable impact on the statistical jump tail. We conjecture that this wider gap between

fundamentals and asset prices may be accounted for through an extension of the preferences via some form for time-varying risk aversion and/or ambiguity aversion towards extreme downside risk.

The rest of the paper is organized as follows. Section 2 describes our data and estimation method. Section 3 presents the results from fitting the option panel using a general version of the standard two-factor stochastic volatility asset pricing model along with our new extension. In Section 4 we exploit the estimation results from Section 3 to study the risk premium dynamics and its implication for return and variance predictability. Section 5 concludes. In a Supplementary Appendix we report estimation results on subsamples as well as for various alternative specifications for the risk-neutral dynamics of the underlying index. The Supplementary Appendix also contains diagnostic tests related to the parametric fit for the option surface.

2 Data and Estimation Methodology

We start by briefly describing our option and underlying asset price data as well as the parametric estimation method we employ for fitting the option panel.

2.1 Data

We use European style S&P 500 equity-index (SPX) options traded at the CBOE. We exploit the closing bid and ask prices reported by OptionMetrics, applying standard filters and discarding all in-the-money options, options with time-to-maturity of less than 7 days, as well as options with zero bid prices. For all remaining options, we compute the mid bid-ask Black-Scholes implied volatility. The data spans January 1, 1996, till July 21, 2010. Following earlier empirical work, e.g., Bates (2000) and Broadie et al. (2009), we sample every Wednesday.[1] The sample includes 760 trading days, and the estimation is based on an average of 234 bid-ask quotes per day. The nonparametric estimate of volatility needed for the penalization of the objective function below is constructed from one-minute high-frequency data on the S&P 500 futures covering the time span of the options. The same data set is used to construct measures of volatility and jump risks for the predictive regressions in Section 4. Finally, we also employ the returns on the SPY ETF traded on the NYSE, which tracks the S&P 500 index portfolio, and the 3 month T-bill rate to proxy for the risk-free rate, when implementing these predictive regressions.

[1] Due to extreme violations of no-arbitrage-conditions, we replaced October 8, 2008, with October 6, 2008.

2.2 Estimation

To define our estimator, we first introduce some notation. We denote European-style out-of-the-money option prices for the asset X at time t by $O_{t,k,\tau}$. Assuming frictionless trading in the options market, the option prices at time t are given as,

$$O_{t,k,\tau} \;=\; \begin{cases} \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} r_s\, ds}\, (X_{t+\tau} - K)^+ \right], & \text{if} \quad K > F_{t,t+\tau}\ , \\ \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{t+\tau} r_s\, ds}\, (K - X_{t+\tau})^+ \right], & \text{if} \quad K \leq F_{t,t+\tau}\ , \end{cases} \tag{2.1}$$

where τ is the tenor of the option, K is the strike price, $F_{t,t+\tau}$ is the futures price for the underlying asset at time t referring to date $t+\tau$, for $\tau > 0$, $k = \ln(K/F_{t,t+\tau})$ is the log-moneyness, and r_t is the instantaneous risk-free interest rate. Finally, we denote the Black-Scholes implied volatility corresponding to the option price $O_{t,k,\tau}$ by $\kappa_{t,k,\tau}$. This merely represents an alternative notational convention, as the Black-Scholes implied volatility is a strictly monotone transformation of the ratio $\frac{e^{r_{t,t+\tau}} O_{t,k,\tau}}{F_{t,t+\tau}}$, where $r_{t,t+\tau}$ denotes the risk-free rate over the period $[t, t+\tau]$.

As for virtually all models used in option pricing, we specify the risk-neutral dynamics of X_t (as well as that of r_t and δ_t) via a Markov state vector and denote with $\mathbf{S}_t$ the part of the latter which determines the conditional distribution of the increments of $\log(X_t)$. In this setting, the model-implied option price at time t (quoted in units of Black-Scholes implied volatility) is given by $\kappa(k, \tau, \mathbf{S}_t, \theta)$, where θ denotes the parameter vector of the model.

Allowing for observation error in the option prices, we denote the average of the observed bid and ask quotes (expressed Black Scholes implied volatility units) by $\overline{\kappa}_{t,k,\tau}$, while the number of intraday returns used each day is given by n. We estimate the period-by-period realization of the state vector as well as the model parameter vector by a penalized weighted least squares procedure,

$$\left(\{\widehat{\mathbf{S}}_t^n\}_{t=1,\ldots,T},\, \widehat{\theta}^n \right) = \underset{\{\mathbf{Z}_t\}_{t=1,\ldots,T},\, \theta \in \Theta}{\operatorname{argmin}} \sum_{t=1}^{T} \left\{ \frac{1}{N_t} \sum_{j=1}^{N_t} \frac{\left(\overline{\kappa}_{t,k_j,\tau_j} - \kappa(k_j, \tau_j, \mathbf{Z}_t, \theta) \right)^2}{\widehat{V}_t^n} \;+\; \zeta_n \frac{\left(\widehat{V}_t^{(n,m_n)} - \xi(\mathbf{Z}_t) \right)^2}{\widehat{V}_t^n} \right\}, \tag{2.2}$$

for a deterministic sequence of nonnegative numbers $\{\zeta_n\}$ (decreasing asymptotically to zero, as n diverges) and $\xi(\cdot)$ being the model-based map from the state vector to the spot diffusive volatility (this map may depend on the parameter vector). $\widehat{V}_t^n \equiv \widehat{V}_t^{(n,n)}$ and $\widehat{V}_t^{(n,m_n)}$ are nonparametric estimators of the diffusive return variation constructed from the intraday record of the log-futures price of the underlying asset, f, as follows,

$$\widehat{V}_t^{(n,m_n)} \;=\; \frac{n}{m_n} \sum_{i=tn-m_n+1}^{tn} (\Delta_i^n f)^2\, 1_{\{|\Delta_i^n f| \leq \alpha n^{-\varpi}\}}, \qquad \Delta_i^n f = f_{\frac{i}{n}} - f_{\frac{i-1}{n}}, \tag{2.3}$$

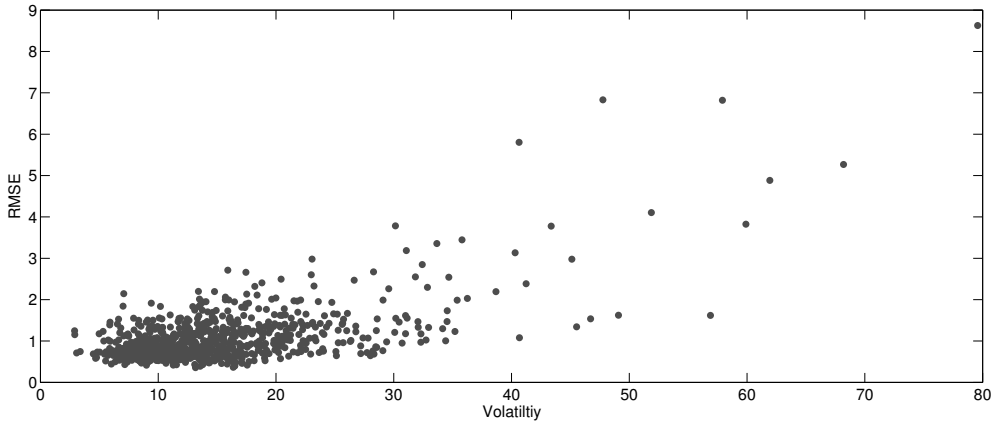

Figure 1: **Option Pricing Root-Mean-Squared Errors as a Function of Volatility.** The data are obtained from the three-factor model estimated by Andersen et al. (2012).

where $\alpha > 0$, $\varpi \in (0, 1/2)$, and m_n denotes a deterministic sequence with either $m_n/n \to 0$ or $m_n = n$. For $m_n = n$, $\widehat{V}_t^{(n,m_n)}$ is the truncated variation (see, e.g., Mancini (2009)), which provides a consistent estimate of the integrated diffusive variance over $[t-1, t]$. For $m_n/n \to 0$, $\widehat{V}_t^{(n,m_n)}$ is a consistent estimator of the spot variance at t and corresponds to the truncated variation computed over an (asymptotically shrinking) fraction of the day just prior to the option quote. In our implementation, we sample every minute over a 6.75 hours trading day, excluding the initial five minutes, resulting in $n = 400$. We employ $m_n = 300$ for $\widehat{V}_t^{(n,m_n)}$. Finally, the tuning parameters α and ϖ are determined in a data-driven fashion, as in Bollerslev and Todorov (2011b).

The estimator in (2.2) is based on minimizing the weighted mean squared error in fitting the panel of observed option implied volatilities, with a penalization term that reflects how much the model-implied spot volatility deviates from a model-free volatility estimate. The presence of $\widehat{V}_t^{(n,m_n)}$ in the objective function serves as a regularization device that helps identify the parameter vector by penalizing values that imply "unreasonable" volatility levels. Moreover, the standardization by $\widehat{V}_t^n$ allows us to weigh option observations on high and low volatility days differently to improve efficiency by accounting for the fact that the option error generally increases with market volatility. This is the point at which our estimation procedure deviates from the empirical analysis in Andersen et al. (2012). The rationale for our weighting scheme is apparent from Figure 1. It demonstrates that the option pricing errors are roughly linear in market volatility, motivating our scaling of the errors by the level of volatility.

Andersen et al. (2012) develops the limiting distribution for the estimator based on an in-fill asymptotic scheme. The estimator is consistent for the parameter vector and the period-by-period

realization of the (latent) state vector, while the distributional convergence to a mixed Gaussian limit is stable, enabling feasible inference and formal diagnostic tests. Beyond details on the formal estimation, inference and testing procedures associated with this approach, Andersen et al. (2012) also document encouraging finite sample performance through a large scale simulation study.

In practice, the optimization in equation (2.2) is most conveniently performed iteratively, exploiting two simple steps. For a given parameter vector, we first determine the optimal value of the state vector on each day of the sample. Second, for the concentrated objective function, we perform MCMC based estimation of the parameter vector with a chain of length $10,000$. For this new θ estimate, we then move back to step one and iterate the procedure.

3 Parametric Modeling of the Option Panel

We start our analysis by fitting the option panel using the data and estimation method described in the previous section. The objective is to capture the dynamics of the option surface succinctly through a parametric stochastic volatility model which may subsequently serve as a proper basis for the analysis of the equity and variance risk premia. A pertinent issue is the number of latent state variables required to provide a suitable fit. The vast majority of empirical option pricing studies employs a single stochastic volatility factor, but the literature on estimating the return dynamics under the physical measure points to a minimum of two factors. Consequently, we start by exploring a couple of alternative one-factor volatility models and, guided by shortcomings in their fit to the option panel dynamics, we progress towards richer specifications.

3.1 Alternative Representations for the Risk-Neutral Asset Price Dynamics

3.1.1 One-Factor Jump-Diffusive Specifications

The popular one-factor volatility model of Duffie et al. (2000), also known as the double-jump model, specifies the risk-neutral dynamics for the price process X as,

$$\begin{aligned} \frac{dX_t}{X_{t-}} &= (r_t - \delta_t)\,dt + \sqrt{V_{1,t}}\,dW_{1,t}^{\mathbb{Q}} + \int_{\mathbb{R}^2} (e^x - 1)\,\widetilde{\mu}^{\mathbb{Q}}(dt, dx, dy), \\ dV_{1,t} &= \kappa_1\,(\overline{v}_1 - V_{1,t})\,dt + \sigma_1\,\sqrt{V_{1,t}}\,dB_{1,t}^{\mathbb{Q}} + \int_{\mathbb{R}^2} y\mu(dt, dx, dy), \end{aligned} \tag{3.1}$$

where $(W_{1,t}^{\mathbb{Q}}, B_{1,t}^{\mathbb{Q}})$ is a two-dimensional Brownian motion with corr $\left(W_{1,t}^{\mathbb{Q}}, B_{1,t}^{\mathbb{Q}}\right) = \rho_1$, while μ is an integer-valued random measure on $\mathbb{R}_+ \times \mathbb{R}^2$ counting the jumps in the price, X, as well as the state vector, $V_{1,t}$, with compensator (under the risk-neutral measure) $dt \otimes \nu_t^{\mathbb{Q}}(dx, dy)$ and $\widetilde{\mu}^{\mathbb{Q}}(dt, dx, dy) = \mu(dt, dx, dy) - dt \otimes \nu_t^{\mathbb{Q}}(dx, dy)$ is the associated martingale measure.

Moreover, the risk-neutral compensator for the jump measure takes the form,

$$\nu_t^{\mathbb{Q}}(dx,dy) \;=\; \left\{ (c_0 + c_1 V_{1,t-}) \; \frac{e^{-\frac{(x-\mu_x-\rho_j y)^2}{2\sigma_x^2}}}{\sqrt{2\pi}\sigma_x} \; \frac{e^{-y/\mu_v}}{\mu_v} \, 1_{\{y>0\}} \right\} dx \otimes dy.$$

This representation implies that the price and volatility processes jump simultaneously. The return jumps are normally distributed with mean μ_X and variance σ_X^2, while the volatility jumps are exponentially distributed with parameter μ_v and constrained to be positive. The correlation in the jump size of returns and volatility is governed by ρ_j. Finally, the jump intensity varies with the level of volatility, and is controlled by the coefficients c_0 and c_1.[2]

Most existing empirical option pricing studies exploit a constrained version of the above model. For example, if we rule out time-variation in the jump intensity, $c_1 = 0$, we obtain the specification used by Broadie et al. (2009). If we rule out volatility jumps, we obtain a model similar to those estimated by Bates (2000) and Pan (2002).[3]

The specification of the risk-neutral measure, $\mathbb{Q}$, determines how the distinct sources of randomness (risk) in the return process are compensated, or priced, in the model. The above setting has some notable features in this regard. First, the price as well as the volatility trajectories display discontinuities. There is strong empirical evidence for both, including possible dependence between them. Second, the jump risk (under the risk-neutral measure) is time-varying. Strong time-variation of the jump risk, under both the statistical and risk-neutral measure, has been documented by nonparametric means in the recent literature. Third, nonetheless, the specification is restrictive in the sense that it implies a single (stochastic volatility) state variable spans the dynamics of the entire option surface across time.

Table 1 provides the parameter estimates for the model, using the data and estimation method reviewed in Section 2. Several features of the risk-neutral measure are noteworthy. First, the volatility factor, V_1, is highly persistent, with shocks to the process having a half-life of about one year. Second, the innovations in price (of diffusive and jump type) are strongly correlated with diffusive volatility ($\rho_1 \approx -0.90$) as well as volatility jumps ($\rho_j \approx -0.62$) inducing a pronounced "leverage" effect. Third, the mean jump size is strongly negative at -9.25% with an average number of jumps per year of 0.63. Fourth, the jump distribution is very wide, indicating that while extremely large negative jumps are possible, so are fairly large positive ones as well. Fifth, the model implies an annualized estimate for the (square-root of the) quadratic return variation of 21.37%. Since

[2]The stationarity conditions for many of our models are standard and relegated to the Supplementary Appendix.

[3]Affine specifications for the asset price also arise naturally within equilibrium models, in which case the latent factor may be given a structural interpretation and components of the risk in the asset price process will be related to underlying fundamentals in specific ways; examples include Drechsler and Yaron (2011) and Drechsler (2013).

the average realized volatility over the sample is about 17.9%, this implies a large volatility risk premium, as also widely reported in the literature. Finally, most of the jump intensity is driven by the time-varying component. This reflects the fact that short-maturity deep out-of-the-money options become relatively more expensive during periods of high volatility.

Table 1: **Parameter estimates for the one-factor Gaussian jump model**

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	-0.899	0.009	c_0	0.011	0.002	σ_x	0.138	0.001
$\overline{v}_1$	0.028	0.001	c_1	20.517	0.347	μ_v	0.025	0.001
κ_1	0.686	0.017	μ_x	-0.076	0.001	ρ_j	-0.621	0.048
σ_1	0.154	0.003						

Note: **Parameter estimates of the one-factor model (3.1)**. The model is estimated using S&P 500 equity-index option data sampled every Wednesday over the period January 1996-July 2010.

We next explore some powerful diagnostics that help identify scenarios in which the model fails. Andersen et al. (2012) develop Z-scores that reflect the fit to specific parts of the option surface during a given trading period. They are obtained as the normalized statistics, $\mathcal{Z}_{\mathcal{K},t}/\sqrt{\widehat{\text{Avar}}\left(\mathcal{Z}_{\mathcal{K},t}\right)}$, based on the average pricing errors over the relevant part of the surface,

$$\mathcal{Z}_{\mathcal{K},t} = \sum_{j:k_j\in\mathcal{K}} \left(\overline{\kappa}_{t,k_j,\tau^*} - \kappa(k_j,\tau^*,\widehat{\mathbf{S}}_t^n,\widehat{\theta}^n)\right),$$

where $\mathcal{K}$ indicates a specific region of moneyness and τ^* a given tenor. Andersen et al. (2012) show that these Z-scores are asymptotically standard normal under the null of correct model specification.

Figure 2 plots the Z-scores for each day in the sample corresponding to separate parts of the option surface. These regions span two maturity ranges – short-term and long-term options – as well as three categories of moneyness – out-of-the-money (OTM) puts, at-the money (ATM) options, and OTM calls – leading to a total of six regions for each trading day.[4]

The one-factor model provides a reasonable fit to the average level of the short-term OTM put options, although the violations of the 95% confidence band naturally become more pronounced

[4]The region are defined as follows. Short maturity OTM put options: $\tau \leq 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \leq -1$; Short maturity ATM options: $\tau \leq 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \in (-1,1)$; Short maturity OTM call options: $\tau \leq 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \geq 1$; Long maturity OTM put options: $\tau > 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \leq -1$; Long maturity ATM options: $\tau > 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \in (-1,1)$; Long maturity OTM call options: $\tau > 60$ days and $\frac{k}{\sigma_t^{ATM}\sqrt{\tau}} \geq 1$.

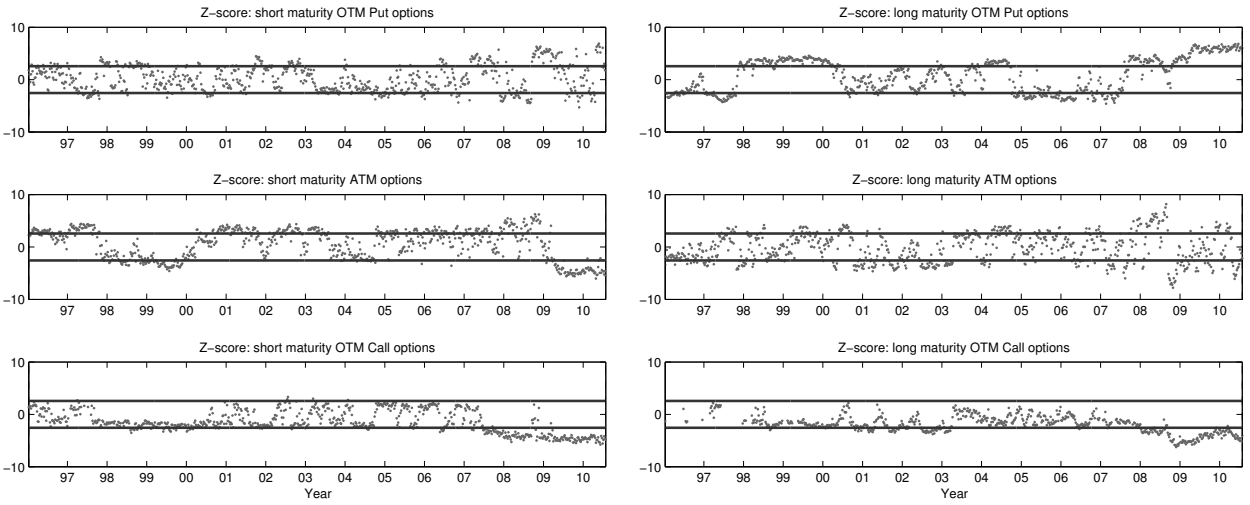

Figure 2: **Z-scores for the one-factor Gaussian jump model (3.1).**

during the 2008-2009 financial crisis. For the two separate ATM categories, the average fit is also sensible, but there are more episodes with systematic mispricing, particularly for the shorter maturities. Finally, we observe extremely large biases for the longer term OTM put options and OTM call options, for which the Z-scores fall outside the confidence bands for long stretches of time. Additional diagnostic tests are reported in a Supplementary Appendix, confirming that the model face severe difficulties in accommodating the full option surface dynamics.

Since the risk-neutral dynamics is a critical input into the modeling of the risk premium dynamics, we seek a more suitable representation. One striking feature of the estimated system is the large negative mean jump size and the low jump intensity. They imply a minimal chance of observing small negative or positive jumps in any given year, but this is at odds with the nonparametric evidence for a gradual decay in the implied jump size for close-to-maturity OTM options in, e.g., Bollerslev and Todorov (2011b).[5] This discrepancy motivates an alternative specification for the jump distribution that designates separate parameters to accommodate the right and left jump tails and allows for smaller discontinuities to occur more frequently than larger jumps.

Hence, we next consider a specification identical to model (3.1) except for the jump distribution. There are two main changes. First, the return jumps now follow a double exponential distribution with separate tail decay parameters, λ_- and λ_-, for the left and right tails. Second, the price-volatility co-jumps are now perfectly dependent, with the squared price jumps impacting the volatility dynamics in a manner reminiscent of a GARCH specification in discrete time. However,

[5]Small jumps are also typically found to be relatively frequent and close to symmetrically distributed by nonparametric estimates based on the underlying asset.

for parsimony and ease of identification, we allow only the negative price jumps to impact the volatility dynamics.[6] The return system takes the form,

$$\begin{aligned} \frac{dX_t}{X_{t-}} &= (r_t - \delta_t)\,dt + \sqrt{V_{1,t}}\,dW^{\mathbb{Q}}_{1,t} + \int_{\mathbb{R}^2} (e^x - 1)\,\widetilde{\mu}^{\mathbb{Q}}(dt, dx), \\ dV_{1,t} &= \kappa_1\,(\overline{v}_1 - V_{1,t})\,dt + \sigma_1\,\sqrt{V_{1,t}}\,dB^{\mathbb{Q}}_{1,t} + \mu_1 \int_{\mathbb{R}^2} x^2\,1_{\{x<0\}}\,\mu(dt, dx), \end{aligned} \tag{3.2}$$

The jump measure μ has a compensator given by $dt \otimes \nu^{\mathbb{Q}}_t(dx)$, where,

$$\nu^{\mathbb{Q}}_t(dx, dy) = \left(c^- 1_{\{x<0\}} \lambda_- e^{-\lambda_- |x|} + c^+ 1_{\{x>0\}} \lambda_+ e^{-\lambda_+ x}\right) dx,$$

$$c^- = c^-_0 + c^-_1 V_{1,t-}, \qquad c^+ = c^+_0 + c^+_1 V_{1,t-}\,.$$

Table 2 reports on the estimation result for model (3.2). As expected, the jump distribution is now dramatically different. The implied mean negative and positive jump sizes are respectively -7.9% and 1.94%, with an average of 1.52 negative and 3.75 positive jumps per year. Thus, jumps are more frequent and smaller – especially for the positive returns – while the negative jumps intensity is roughly proportional to the level of volatility (c^-_0 was insignificant and fixed at zero). We also note the distinctly lower degree of tail decay for the left relative to the right tail. As before, the evidence of co-jumps is strong as the parameter μ_1 is highly significant. Finally, while the remaining parameters also shift relative to model (3.1), the changes are minor so the diffusive component is qualitatively similar to before. Specifically, the mean volatility level has barely changed and the implied half life of a volatility shock remains close to one year.

Table 2: **Parameter estimates for the one-factor Exponential jump model**

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	−0.833	0.006	c^+_0	3.185	0.063	λ_-	12.652	0.071
$\overline{v}_1$	0.018	0.001	c^-_1	55.110	0.780	λ_+	51.616	0.221
κ_1	0.770	0.013	c^+_1	20.567	4.342	μ_1	1.027	0.020
σ_1	0.157	0.002						

Note: **Parameter estimates of the one-factor model (3.2)**. The model is estimated using S&P 500 equity-index option data sampled every Wednesday over the period January 1996-July 2010.

The diagnostic tests presented in the Supplementary Appendix reveal a non-trivial improvement in the overall fit. For example, the root-mean-square error (RMSE) of the model-implied fit to the

[6] Support for this type of price-volatility jump dependence, albeit under $\mathbb{P}$, may be found in Todorov (2011).

observed (implied volatilities of the) option data are now 2.60% compared to 3.14% for model (3.1). Nonetheless, the qualitative fit to the option surface dynamics is otherwise not substantially improved and identical problems to before are evident from the diagnostic tests.

3.1.2 Two-Factor Jump-Diffusive Specifications

The recent literature has documented that one-factor jump-diffusive models are unable to capture the return dynamics for equity indices and multi-factor models are now common. Thus, it seems natural to extend the standard setting for option pricing to a multi-factor setting as well. Our estimation and inference framework can readily accommodate such systems, as demonstrated by the simulation evidence and empirical illustrations in Andersen et al. (2012). We therefore extend the above models to a two-factor stochastic volatility setting. One of these representations mimics the initial specification explored in Andersen et al. (2012), but we now estimate the system more efficiently by weighting the criterion function by current volatility.

In the Supplementary Appendix, we report estimation results and diagnostic tests for the two-factor extension to both one-factor models considered above. Since we again find the double-exponential jump specification to provide the superior fit, we focus on that representation here. The model involves the following risk-neutral dynamics,

$$\begin{aligned} \frac{dX_t}{X_{t-}} &= (r_t - \delta_t)\,dt + \sqrt{V_{1,t}}\,dW^{\mathbb{Q}}_{1,t} + \sqrt{V_{2,t}}\,dW^{\mathbb{Q}}_{2,t} + \int_{\mathbb{R}^2} (e^x - 1)\,\widetilde{\mu}^{\mathbb{Q}}(dt,dx), \\ dV_{1,t} &= \kappa_1\,(\overline{v}_1 - V_{1,t})\,dt + \sigma_1\,\sqrt{V_{1,t}}\,dB^{\mathbb{Q}}_{1,t} + \mu_1 \int_{\mathbb{R}^2} x^2\,1_{\{x<0\}}\,\mu(dt,dx), \\ dV_{2,t} &= \kappa_2\,(\overline{v}_2 - V_{2,t})\,dt + \sigma_2\,\sqrt{V_{2,t}}\,dB^{\mathbb{Q}}_{2,t}. \end{aligned} \tag{3.3}$$

The jump measure μ has a compensator given by $dt \otimes \nu^{\mathbb{Q}}_t(dx)$, where,

$$\nu^{\mathbb{Q}}_t(dx,dy) = \left\{ \left(c^- 1_{\{x<0\}} \lambda_- e^{-\lambda_- |x|} + c^+ 1_{\{x>0\}} \lambda_+ e^{-\lambda_+ x} \right) \right\} dx,$$

$$c^- = c^-_0 + c^-_1 V_{1,t-} + c^-_2 V_{2,t-}, \qquad c^+ = c^+_0 + c^+_1 V_{1,t-} + c^+_2 V_{2,t-}\,.$$

Relative to model (3.2), the volatility is now driven by two independent components: $V_{1,t}$ and $V_{2,t}$. We allow for one of the factors, V_1, to display discontinuities and potentially co-jump with the asset price. Moreover, the jump intensities are now linked differentially to the two factors, allowing for additional flexibility in the characterization of the jump dynamics. In addition, the return innovations may be correlated with the innovations in both volatility factors.

Table 3 provides parameter estimates. It is evident that the two volatility factors take on distinct roles. The factor co-jumping with returns, V_1, is relatively small, transient and variable, while V_2 is on average larger, less volatile and highly persistent. The half life for volatility shocks

Table 3: **Parameter estimates for the two-factor Exponential jump model.**

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	-0.762	0.091	κ_2	0.169	0.014	c_2^+	52.795	4.597
$\overline{v}_1$	0.004	0.000	σ_2	0.129	0.002	λ_-	16.943	0.070
κ_1	12.831	0.227	c_0^+	2.315	0.063	λ_+	51.818	0.220
σ_1	0.247	0.031	c_1^-	69.208	1.856	μ_1	6.262	0.143
ρ_2	-0.945	0.009	c_1^+	13.161	11.356			
$\overline{v}_2$	0.059	0.004	c_2^-	97.656	1.091			

Note: **Parameter estimates of the two-factor model (3.3)**. The model is estimated using S&P 500 equity-index option data sampled every Wednesday over the period January 1996-July 2010.

is 0.05 (a few weeks) and 4.1 years, respectively. Many of the parameters controlling basic features of the return dynamics in model (3.2) are now represented by a pair of coefficients through the loadings on the volatility factors. For example, the diffusive leverage coefficients in model (3.3), ρ_1 and ρ_2, straddle the value in model (3.2), with ρ_1 smaller (in absolute value) and ρ_2 larger than before, thus indicating an overall effect of similar magnitude. The same happens for the mean volatility, with $\overline{v}_1$ being relatively low and $\overline{v}_2$ high, as well as κ_1 being comparatively large and κ_2 small. The various jump parameters now load on new factors and are not readily comparable to before, but the tail decay coefficients, λ_- and λ_+, remain close to the prior point estimates. Hence, overall, the estimated jump distributions have not changed much in response to the introduction of an additional volatility factor. This is corroborated by the implied mean jump intensities of 2.6 and 3.4 jumps per year and mean jump sizes of -5.9% and 1.9%, respectively, for the negative and positive return jumps. The only noticeable change is a minor increase in the incidence of negative jumps and a corresponding small decline in the (absolute) size of these jumps.

The two-factor model represents a substantial improvement in the fit to the option-implied volatility surface, with a RMSE for the model-implied fit of 2.07% compared to 2.60% for model (3.2). This can also be illustrated by visual means through the Z-score diagnostics in Figure 3. It is evident that the main improvement occurs for the ATM options. This is quite intuitive as the diffusive volatility is of primary importance in pricing the ATM options, while the jump features become critical for short maturity OTM options. Hence, the improved fit likely reflects the more flexible dependence structure in volatility, stemming from the inclusion of an additional factor.

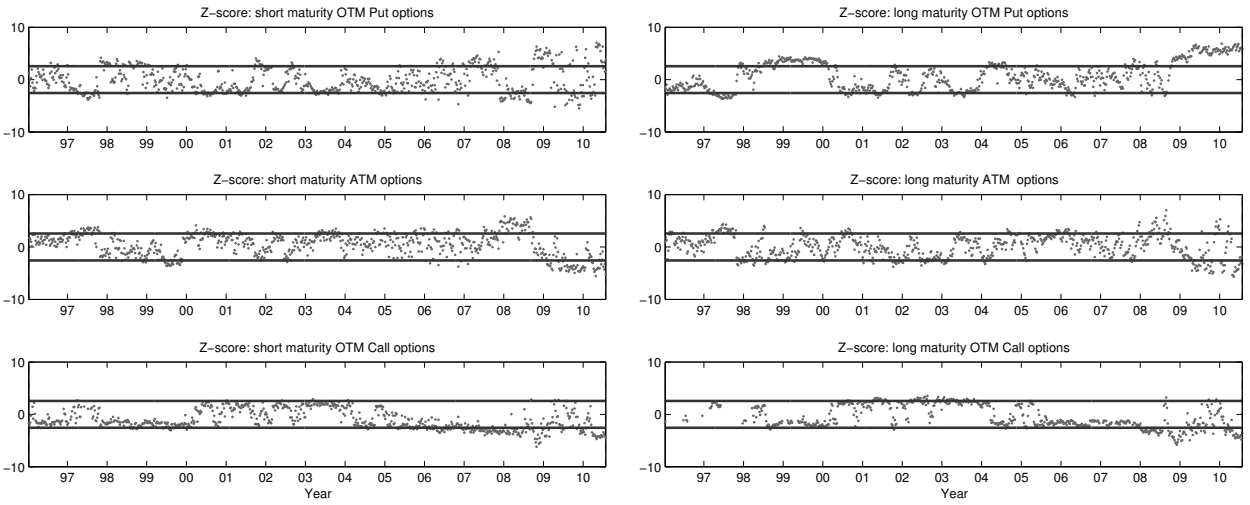

Figure 3: **Z-scores for the two-factor Exponential jump model (3.3).**

This also manifests itself in a marginal improvement in fit for the longer term OTM put options, as can be verified from the model diagnostics in the Supplementary Appendix. However, there fit does not improve much for the remaining option categories. Overall, the model clearly has problems capturing some important aspects of the volatility surface dynamics.

To gauge the shortcomings of model (3.3), we inspect the fit to the option surface across time, as reflected in the Z-scores on Figure 3, more closely. On average, or unconditionally, there are only minor systematic biases, as the points are scattered quite evenly across the confidence band – and beyond. However, there are readily discernible patterns in the Z-scores, indicating specific periods of significant model failure. In particular, the pricing errors are not only highly persistent, but also negatively correlated across option categories. For the crisis periods, 1998-2000 and 2008-2010, the longer term OTM put options are severely underpriced, while the short term ATM options and all OTM call options are overpriced by the model. The reverse happens during the comparatively tranquil market conditions before 1997 and around 2003-2004. Clearly, this suggests there is an inherent conflict in pricing these separate groups of options during specific market conditions. For example, during crisis, the short term OTM puts are expensive relative to the ATM options or, in other words, the volatility smirk is unusually steep. The model can only generate this gap by enhancing the intensity of negative price jumps but, in the current model, they are linked to the state vector representing components of spot volatility. When implied volatility rises to accommodate the valuation of the short term put options, the long maturity OTM put options remain underpriced. This suggest the left tail of the jump distribution displays a persistent type of variation that cannot be spanned by a linear function of the volatility factors. A potential remedy is the introduction of

an additional state variable, allowing the left jump tail to also vary independently of the volatility factors. That is, the time-variation in the pricing of negative jump risk may to some extent evolve separately from volatility, even if the two remain correlated.

In summary, a steep implied volatility curve and a flat term structure of implied volatility for OTM put options during turbulent market conditions appear to be systematic features of the option surface dynamics. These episodes induce severe pricing violations, because the jump intensity is an affine function of the volatility state. More formally, near maturity ATM Black-Scholes implied volatility approximates the current (diffusive) volatility well, see, e.g., Durrleman (2008), while short-maturity OTM put options are almost proportional to the jump intensity, see, e.g., Carr and Wu (2003). Therefore, the volatility level basically fixes the slope of the short-maturity implied volatility curve in the model. These observations motivate the extension below.

3.1.3 The Extended Three-Factor Jump-Diffusive Model

Following Andersen et al. (2012). we extend the two-factor model by stipulating the following risk-neutral equity index dynamics,

$$\begin{aligned}
\frac{dX_t}{X_{t-}} &= (r_t - \delta_t)\,dt + \sqrt{V_{1,t}}\,dW^{\mathbb{Q}}_{1,t} + \sqrt{V_{2,t}}\,dW^{\mathbb{Q}}_{2,t} + \int_{\mathbb{R}^2} (e^x - 1)\,\widetilde{\mu}^{\mathbb{Q}}(dt,dx,dy),\\
dV_{1,t} &= \kappa_1\,(\overline{v}_1 - V_{1,t})\,dt + \sigma_1\,\sqrt{V_{1,t}}\,dB^{\mathbb{Q}}_{1,t} + \mu_1 \int_{\mathbb{R}^2} x^2\,1_{\{x<0\}}\,\mu(dt,dx,dy),\\
dV_{2,t} &= \kappa_2\,(\overline{v}_2 - V_{2,t})\,dt + \sigma_2\,\sqrt{V_{2,t}}\,dB^{\mathbb{Q}}_{2,t},\\
dU_t &= -\kappa_3\,U_t\,dt + \mu_u \int_{\mathbb{R}^2} \left[(1-\rho_3)\,x^2\,1_{\{x<0\}} + \rho_3\,y^2\right]\mu(dt,dx,dy).
\end{aligned} \tag{3.4}$$

The jump measure μ has a compensator given by $dt \otimes \nu^{\mathbb{Q}}_t(dx,dy)$, where,

$$\nu^{\mathbb{Q}}_t(dx,dy) = \left\{\left(c^-1_{\{x<0\}}\lambda_- e^{-\lambda_-|x|} + c^+1_{\{x>0\}}\lambda_+ e^{-\lambda_+ x}\right)1_{\{y=0\}} + c^-1_{\{x=0,\ y<0\}}\lambda_- e^{-\lambda_-|y|}\right\}dx\otimes dy,$$

$$c^- = c^-_0 + c^-_1 V_{1,t-} + c^-_2 V_{2,t-} + c^-_3 U_{t-}, \qquad c^+ = c^+_0 + c^+_1 V_{1,t-} + c^+_2 V_{2,t-} + c^+_3 U_{t-}\,.$$

The new specification of the jump dynamics necessitates a modification in notation. As in the prior models, the first term in the jump measure controls the price-volatility co-jumps, but it now also contains a second component governing the independent jumps of the new factor, U. These jumps are either independent from the volatility jumps (in V_1) or proportional to them. Thus, the jump components of U have identical marginal distributions, but differ in their interaction with the remaining jump processes in the model. Finally, the second jump term of μ is defined over the squared negative realizations, which retains the notational symmetry between the jump components of U. We note that this ensures all jumps in U are positive.

Our three-factor model possesses a number of distinctive features. First, we allow only negative price jumps to impact the volatility dynamics, as a more general dependence structure between volatility and price jumps cannot be well identified from the option panel. Second, the negative jumps now have an additional source of time variation, captured by the process U.[7] This feature addresses the shortcomings of the two-factor model (3.3). In particular, U contains an independent source of variation which is driven, in part, by the negative price jumps, enabling the jump intensity to respond vigorously during crisis periods. In addition, we allow for interaction between jumps in U and X through the ρ_3 coefficient. As such, the model accommodates both perfect dependence ($\rho_3 = 0$) and full independence ($\rho_3 = 1$) between the jump risks of V_1 and U. Moreover, these state variables, governing important features of the option panel dynamics, are related through the time-variation in the jump intensity. This feature induces "cross self-excitation" in which jumps in V_1 enhance the probability of future jumps in U, and vice versa.[8]

Table 4: **Parameter Estimates for the extended three factor model**

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	-0.913	0.028	σ_2	0.110	0.006	c_2^-	0.913	3.730
$\overline{v}_1$	0.007	0.000	μ_u	1.756	0.647	c_2^+	14.269	4.986
κ_1	8.325	0.167	κ_3	0.522	0.080	c_3^-	19.836	5.460
σ_1	0.323	0.014	ρ_3	0.117	0.614	λ_-	21.157	0.240
ρ_2	-0.945	0.036	c_0^+	0.723	0.079	λ_+	48.365	2.053
$\overline{v}_2$	0.016	0.001	c_1^-	34.592	1.931	μ_1	11.602	0.262
κ_2	0.480	0.033	c_1^+	88.178	14.711			

Note: **Parameter estimates of the three-factor model (3.4)**. The model is estimated using S&P 500 equity-index option data sampled every Wednesday over the period January 1996-July 2010.

The parameter estimates for model (3.4) are reported in Table 4. The basic features of the volatility factors, V_1 and V_2, are similar to those obtained via model (3.3). Nonetheless, the size of

[7] This feature can also be incorporated into the dynamics for the intensity of the positive price jumps, but since the relevant parameter is not significant, we impose the constraint $c_3^+ = 0$ in our final specification.

[8] We stress that the model (3.4) still belongs to the affine family covered by Duffie et al. (2003) and, as shown in Andersen et al. (2012), the following parameter constraints ensure covariance stationarity of the latent factors,

$$\kappa_1 > \frac{2c_1^-\mu_1}{\lambda_-^2}, \quad \text{and} \quad \kappa_3 > \frac{2c_3^-\kappa_1\mu_u}{\kappa_1\lambda_-^2 - 2c_1^-\mu_1}, \quad \text{and} \quad \kappa_2 < 0, \quad \text{and} \quad \sigma_i^2 \leq 2\kappa_i\overline{v}_i, \quad i = 1, 2.$$

V_2 has declined significantly relative to V_1. Moreover, V_2 is less persistent than before with a half life of about 18 months versus one month for V_1. In addition, the volatility factors exhibit a stronger negative correlation with the diffusive price innovations. While most of the jump parameters are not readily comparable to those in the preceding models, we note that the tail decay parameters only change slightly and the overall features are quite similar to what we established with model (3.3). For example, the negative and positive jump intensities are now about 3.4 and 2.6 per year, while the jump sizes are about -4.7% and 2.1%, respectively, representing only minor offsetting changes in the jump intensities and sizes. Likewise, we again find that the jump intensity is driven almost exclusively by the time-varying component.[9] We also observe that the c_1^- and c_1^+ as well as the c_2^- and c_2^+ coefficients are pairwise statistically distinct, indicating a different degree of time-variation in the left and right jump tail, consistent with the nonparametric evidence in Bollerslev and Todorov (2011b). Further, the state variable U is highly persistent and, as discussed in detail later, is the primary determinant behind the time variation of the left jump tail. Finally, the parameter ρ_3 cannot be estimated with precision, implying that the degree of independent jump variation in U is hard to identify from the option panel alone.

The extended three-factor model (3.4) improves the fit to the option surface substantially relative to model (3.3). The RMSE for the option-implied volatilities drops by close to 15% and now equals 1.75%. Even more strikingly, the Z-scores in Figure 4 portray a set of statistics that largely reside within the 95% confidence bands. The tabulation of option pricing violations at the 1% level across the six strike-maturity categories in the Supplementary Appendix corroborates the point. The week-by-week conditional fit for model (3.3) versus (3.4) yields violation frequencies ranging from 20% to 39% versus 4% to 15%. That is, every option category is more suitably priced in the three-factor model than is the case for the best fit obtained across all six categories in the two-factor setting. Thus, although evidence of persistent mispricing remains, for example for the OTM put and call options during crisis periods, the violations are smaller and more short-lived, and specifically, the ATM options are priced accurately. We conclude that model (3.4) captures the critical dynamic features reasonably well through the stipulated law of motion for the three state variables. Consequently, the residual mispricing should exert less of an impact on our subsequent analysis relative to the more standard specifications explored above.[10]

[9] As before, the estimates of c_0^- were insignificant and fixed at zero.

[10] In principle, some of the remaining valuation problems can be rectified via a time-varying jump tail index, i.e., allowing the parameter λ_- to vary, see, e.g., the nonparametric evidence in Bollerslev and Todorov (2013). However, time variation in the tail index takes us outside the affine setting, implying that we lose the analytical tractability in pricing the option surface which is critical for our current implementation procedures.

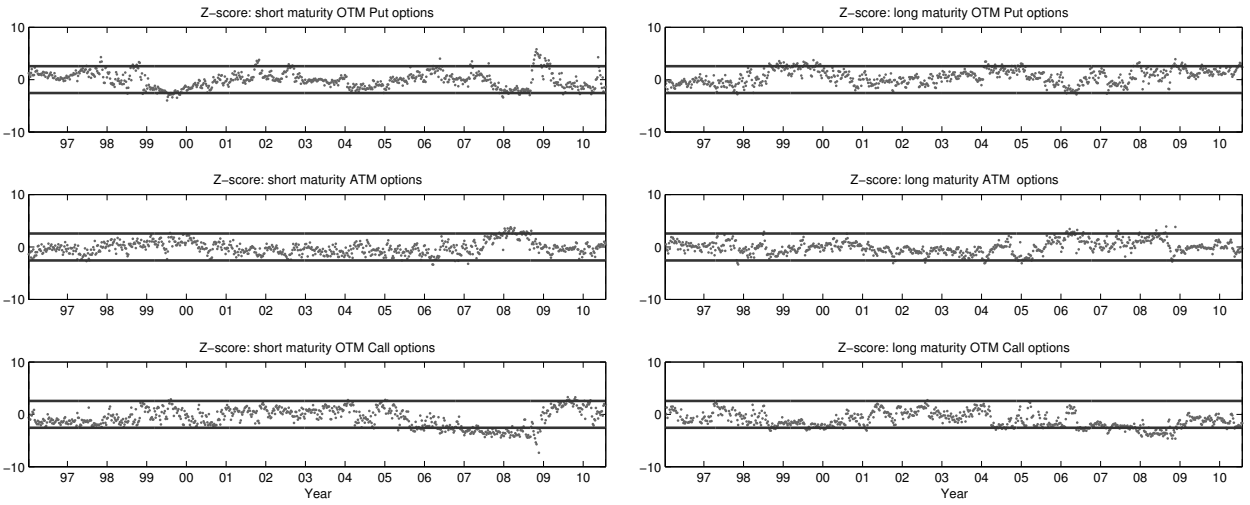

Figure 4: **Z-Scores for the extended three-factor model.**

3.2 Properties of Alternative Models for the Risk-Neutral Price Dynamics

We have established that model (3.4) provides a much improved statistical fit to the option surface relative to more standard representations of the risk-neutral asset price dynamics. However, it is less evident whether the superior fit translates into a better characterization of the key dynamic factors that determine the future evolution of, and associated risks related to, the implied volatility surface as well as the corresponding risk premiums. The remainder of this section is dedicated to illustrating the critical differences in the implied volatility surface dynamics across the alternative models. The question concerning the risk premiums is addressed in the following section.

3.2.1 State Vector Dynamics, Volatility Surface Dynamics, and Jump Intensities

In the estimation step for the three-factor model, we extract estimates for the realizations of the state vector $(V_1\,, V_2\,, U)$ for each week in the sample. Since the state vector, along with the point estimates for the various model parameters, governs the (diffusive) volatility as well as the jump probabilities in the model, they directly imply a path for volatility and for both positive and negative jump intensities over time. Figure 5 displays the implied time series for volatility and the negative jump intensity over our sample period.

There are many fascinating aspects to the plots in Figure 5. First, we see spikes in market volatility and – even more strikingly – in the negative jump intensity around well-known crises. Second, the negative jump intensity portrays a very different picture than the (diffusive) volatility. For example, the peaks of the (negative) jump intensity in the 1997 and 1998 crises as well as the sovereign debt crisis in Europe match or exceed those observed from 2000 through 2003, yet the

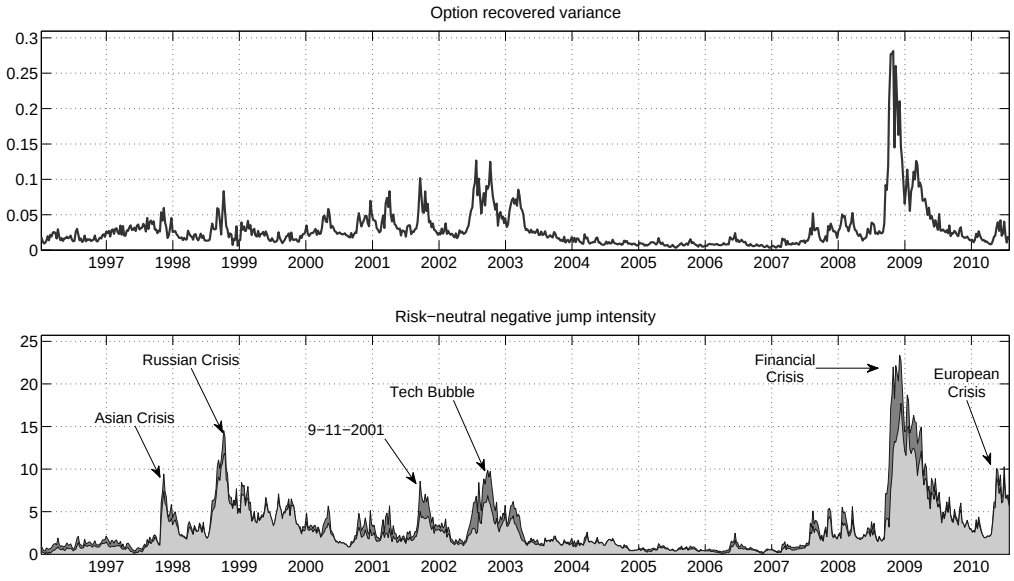

Figure 5: **Three-factor model option-implied spot return variance and negative jump intensities.** The top panel displays the implied annualized spot diffusive variance, while the top line of the bottom panel provides the total negative jump intensity, i.e., $c_0^- + c_1^- V_{1,t-} + c_2^- V_{2,t-} + c_3^- U_{t-}$. The light gray region in the bottom panel indicates the negative jump intensity stemming from U, that is, the $c_3^- U_{t-}$ component.

model-implied volatility is substantially higher in the latter period than during the 1997-1998 episodes or the European crisis. Hence, the alternating shapes of the option surface signal that the episodes represent very different types of exposure to volatility and (negative) jump risks. Third, not surprisingly, the financial crisis stands out as the most dramatic period. It is also evident that jump intensities mean-revert less quickly following the financial crisis. This pattern is repeated across all the episodes identified above in which the jump intensity rises strongly relative to the volatility. Fourth, focusing on the part of the jump intensity that is accounted for by the new jump factor U, we see a similar pattern. For the Asian, Russian and European crises, the elevated jump intensity is almost exclusively driven by U, while this is not the case for the volatile episodes surrounding 9-11 and the bursting of the internet bubble in 2002. Likewise, the initial spike in the jump intensity during the financial crisis was not caused solely by the U factor, but the latter was clearly the determinant behind the very slow reversion thereafter. Finally, we note that the overall pattern for both volatility and jump intensities qualitatively matches the corresponding figure in Andersen et al. (2012), based on the identical model, but estimated via a pure RMSE criterion for the option-implied volatility. The current model is estimated by the same criterion, but applied to

the weighted option-implied volatility errors. The weighting allows the model to accommodate the shape of the option surface equally across different volatility states, while the pure (unweighted) RMSE criterion forces the model to fit the surface as closely as possible during periods with extreme option price fluctuations. This effect is evident here, as the model-implied variance estimates in Andersen et al. (2012) peaks above 0.35 during the financial crisis and also features other values around 0.30. In contrast, for the current estimation approach, the volatility estimates never exceed 0.28. The discrepancies between the two sets of underlying parameter estimates are discussed further in the following subsection.

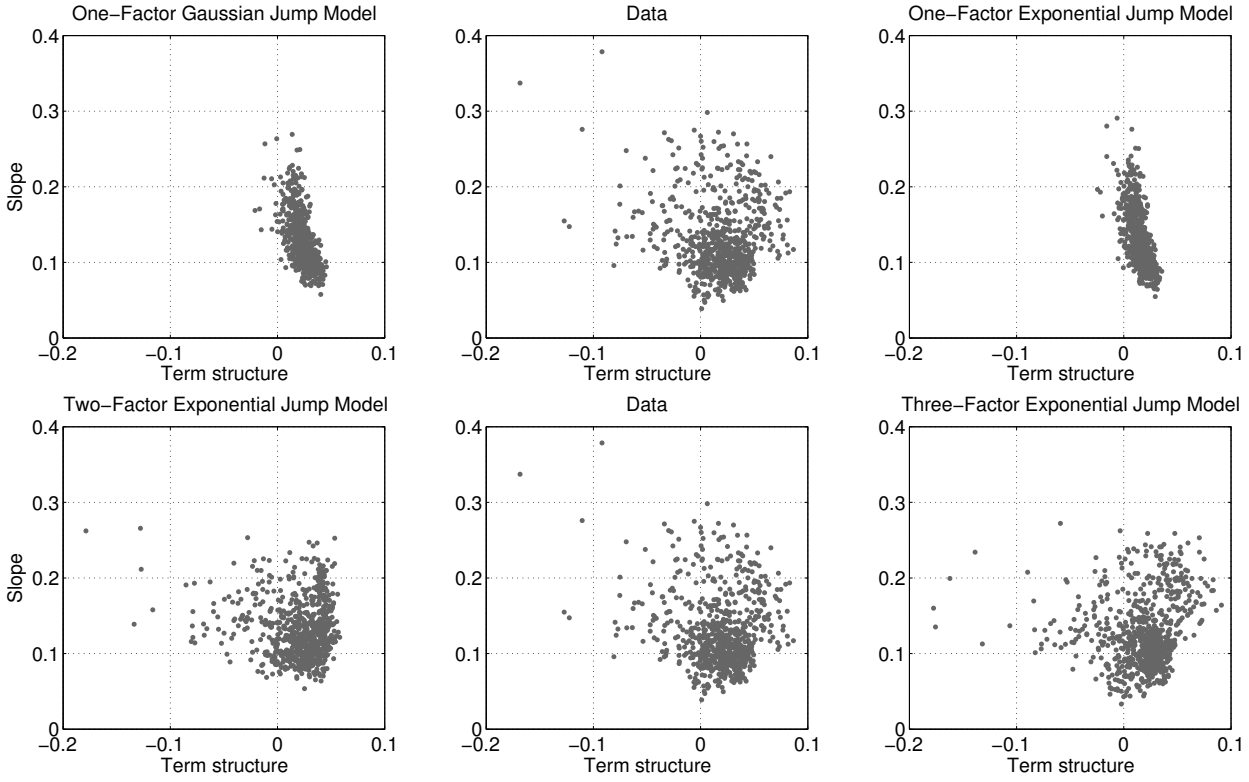

Figure 6: **Scatter plot for the skew versus term structure of implied volatility.** The top row displays the actual observations from the S&P 500 options across our full sample in the center panel, while the corresponding model-implied term structure-skew combinations for model (3.1) and (3.1) are displayed in the left and right panels, respectively. Likewise, the bottom row contains the actual observations (repeated) in the center panel, while the model-implied term structure-skew combinations for models (3.4) and (3.3) across the sample are provided on the left and right. The skew is defined as the difference of the average BSIV for options within the short maturity OTM and ATM categories. The term structure is given by the difference of the average BSIV for the options in the long and short maturity ATM categories.

Figure 5 summarizes key time series implications of model (3.4). In order to understand whether the model represents a major deviation from the one- and two-factor models estimated previously,

we compare the implied properties of the resulting systems along critical dimensions. Figure 6 plots pairwise model-implied slopes for the implied volatility term structure and the volatility skew for all dates in our sample and contrasts them with corresponding values obtained from actual option prices. The first row demonstrates, strikingly, how the one-factor jump diffusions fail to generate the diverse combinations of volatility term structure and smirk slopes observed in the data. For both models, the points are clustered around a steep negatively sloped line. This contrasts dramatically with the combinations obtained from the options which are much more widely dispersed. Specifically, the term structure of implied volatility can be sharply upward or downward sloping when the implied volatility skew is steep. The one-factor models simply cannot generate this feature. The second row shows that the two-factor model provides a major step forward in this respect. However, the scatter plot is too asymmetric relative to the actual data. In particular, model (3.3) fails to generate scenarios with a steep term structure and relatively low implied volatility skew. This feature is accommodated by model (3.4), illustrating the added flexibility afforded by the new, partially independent, jump factor.

To further clarify the distinct role played by the different factors in driving the dynamics of the option panel in model (3.4), Figure 7 depicts the shift in the short and long maturity implied volatilities across the strike prices due to a shock in each of the three state variables, perturbing them from their mean sample values. To provide transparency regarding the critical modeling of the left tail dynamics, we calibrate the shocks so that the shift in V_1 and U have the identical effect on the intensity of negative price jumps. In practical terms, that implies that the shocks to V_1 and U have the same impact on the prices of short-maturity deep OTM put options.[11] Since V_2 has a minimal impact on the left jump tail, we calibrate the shock to this state variable to ensure that its effect on the ATM implied volatility falls within the range of the changes brought about by the shocks to V_1 and U. It is evident from Figure 7 that U, unlike V_1, has a negligible impact on the pricing of short term ATM options. Instead, it primarily influences the pricing of short term deep OTM put options. This is, of course, a reflection of the fact that U is not part of the diffusive volatility. In contrast, shocks to V_1 elevates the entire short-maturity implied volatility curve, with the effect being the largest for the ATM options. In other words, positive shocks to V_1 are ineffective in steepening the implied volatility curve. Turning towards longer maturity options, we again see the shock to V_1 having a smaller impact on the OTM puts than the ATM options, while a shock to U has a similar impact on the entire implied volatility function. In particular, the strong persistence of the U process ensures that it has a much stronger influence on the pricing

[11] As the time-to-maturity of OTM put shrinks towards zero, the option price scaled by time-to-maturity multiplied by the current spot price converges to $\int_{\mathbb{R}}(e^k - e^x)^+ \nu_t^{\mathbb{Q}}(dx, dy)$, see e.g., Carr and Wu (2003).

of the long term options that V_1. Finally, we note that the influence of V_2 is mostly reserved for the ATM and OTM call options for both maturities. In sum, by highlighting the separate roles of the factors in driving the dynamics of the option surface, Figure 7 underscores the critical function of U in accommodating time variation in both the steepness of the volatility smirk and the term structure of implied volatility.

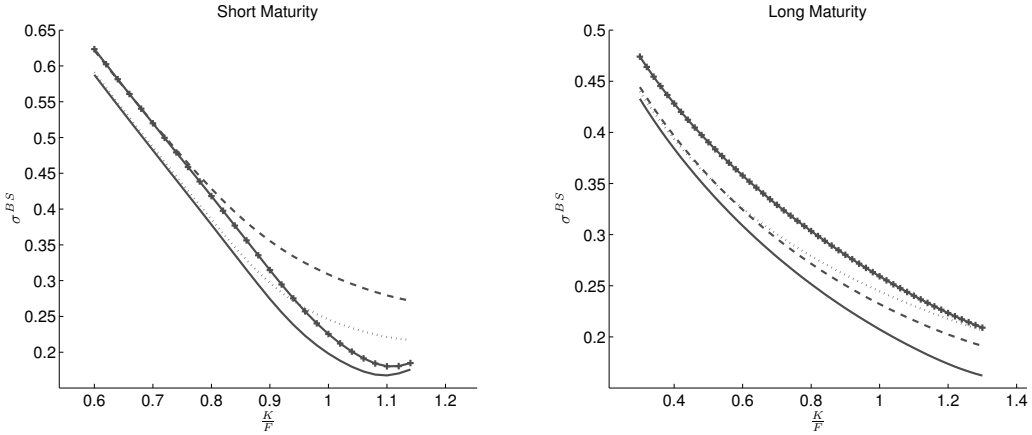

Figure 7: Sensitivity of option surface to changes in state variables in the three-factor model. Parameters are set equal to their estimates reported in Table 4. The solid line corresponds to the implied volatility for values of the three state variables equal to their sample means. All other lines correspond to changing one of the state variables from its sample mean value. The dashed line corresponds to changing V_1, the dotted line to changing V_2 and the x-line to changing U. The change in V_1 is three times its mean value; the change in U is such that the effect on the risk-neutral jump intensity of the negative jumps is the same as that due to changing V_1; the change in V_2 is twice its mean value. Short maturity is one-month and long maturity is one year.

The enhanced flexibility in capturing joint movements in the volatility term structure and skew, attained through the inclusion of the U factor, suggests that a key attribute of model (3.4) is the decoupling of (negative) jump risk from volatility risk. One direct way to assess this feature is to compute the fraction of the overall model-implied conditionally expected (risk-neutral) return variation that is attributable to negative jumps over the sample. Obviously, as we progress from the one- to the two-factor models, we loosen the links between these risk factors, but only the introduction of U allows for a significant degree of independent jump variation. Thus, if the realized option surface dynamics can be better reconciled through this type of mechanism, we would expect to see a significant degree of discrepancy across the models along this dimension. Figure 8 provides this comparison. The implied (negative) jump variation is, indeed, dramatically distinct in our three-factor model. While the fraction of the quadratic return variation accounted for by negative

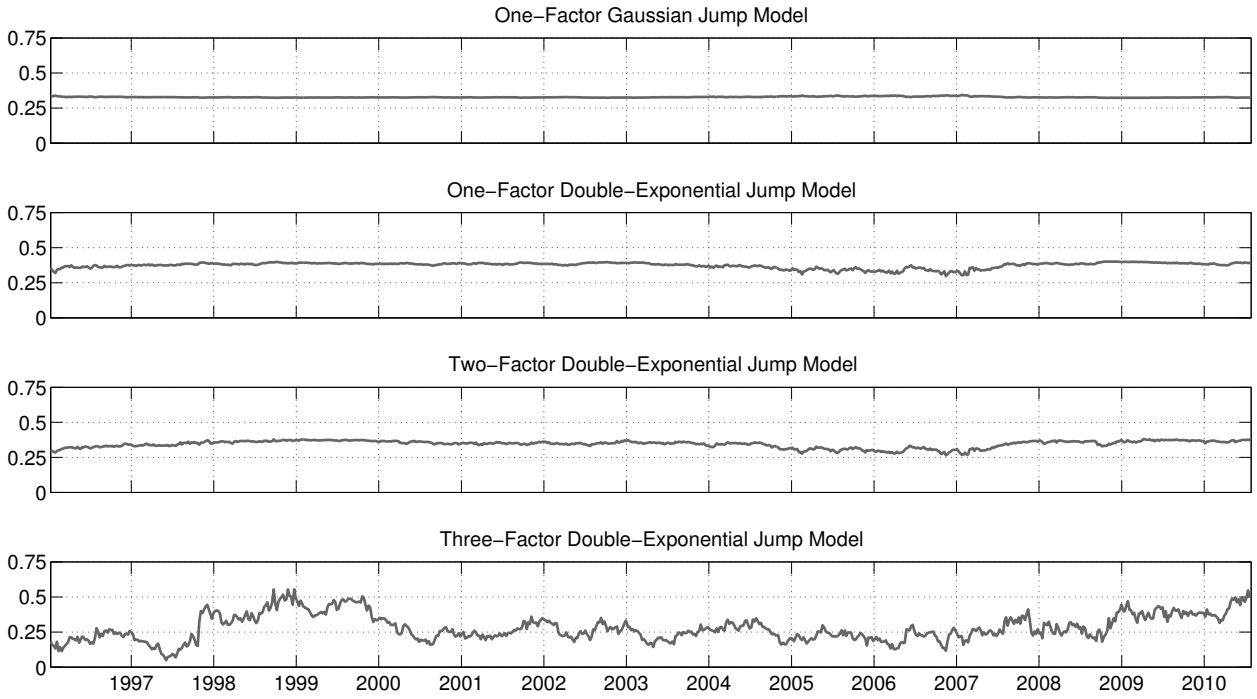

Figure 8: **Ratio of model-implied negative jump variation to total return variation.** Alternative model-implied ratios of negative jump return variation versus total quadratic return variation across the full sample.

jumps is around 30% for all models, it is essentially flat for the one-factor Gaussian jump model, it displays only very minor fluctuations in the one- and two-factor exponential jump models, but it varies significantly, albeit quite smoothly, from under 10% to over 50% for the flexible three-factor model. In particular, the negative jump variation is elevated in the latter model from the onset of the Asian crisis in 1997 up to year 2000, and then again in the aftermath of the financial crisis in 2009 through the end of the sample. As noted before, the initial phase of the financial crisis is not characterized by any large movement in the jump risk relative to volatility risk – both are extremely elevated at that point. The relative importance of the jump risk manifests itself clearly only after volatility starts subsiding. Likewise, the events surrounding the dissolution of LTCM in 1998 and the European sovereign debt crisis in 2010 feature dramatic relative increases in the downside tail risk.

3.2.2 Robustness of Estimation Results and Inference

We next explore the robustness of our empirical findings in a number of directions. First, as indicated previously, our three-factor model (3.4) has been estimated already in Andersen et al. (2012), but our weighting of the criterion function in the current implementation should improve the efficiency and finite sample robustness of the inference. We already noted in Section 3.2.1 that the qualitative features of the extracted factors over the sample are similar across the alternative estimation procedures. The most striking difference in point estimates for the parameters concerns the second volatility factor, V_2. The estimates obtained via the unweighted criterion in Andersen et al. (2012) imply that this factor is substantially larger, more persistent, has more volatile innovations and displays less negative correlation with the returns than is the case for the present estimates. This discrepancy is likely to impact the quality of fit. In fact, the formal diagnostics, reported in the Supplementary Appendix, corroborate this conjecture. The unweighted RMSE criterion, by construction, minimizes the RMSE, which equals 1.61% in Andersen et al. (2012) compared to 1.75% in the current implementation. However, the overall fit to the option surface is clearly superior for the current estimates. The pricing violations at the 1% significance level across our six separate option categories are uniformly higher for the estimates obtained via the unweighted criterion function, ranging from 18% to 37% in Andersen et al. (2012), than reported here, where they fall between 4% and 15%. This uniform improvement is also evident visually when comparing Figure 4 to the corresponding Z-score statistics in Andersen et al. (2012). This suggests that outliers are unduly influential in the inference based on the unweighted objective function relative to the weighted RMSE criterion used in the current work.

Second, one may wonder if the improved performance simply reflects the advantages afforded by a three-factor model relative to less richly parameterized one- and two-factor specifications. Thus, we compare the current model to estimation results obtained from a more traditional representation where the third factor is also a volatility factor in line with those in model (3.3). The exact model specification, parameter estimates and diagnostic tests are presented in the Supplementary Appendix. The overall fit is only slightly worse than for the present model – with a RMSE of 1.86% versus 1.75%. However, as above, the model fails to provide a consistent fit to the option surface across time. Although the model is estimated by the identical weighted criterion function, the pricing violations at the 1% level for the six different option categories range from 13% to 45% versus the 4%-15% span for model (3.4). Moreover, every single option category has a significantly lower number of violations across the sample for the current model. Specifically, for the short maturities, going from OTM puts through ATM options to OTM calls, the violations are 21%

versus 8%, 13% versus 4%, and 35% versus 15%. Moreover, the relative fit is even worse for the more traditional model at long maturities. Hence, the introduction of a partially independent (negative) jump factor is critical for accommodating the shifting shape of the option surface. In other words, the option surface dynamics provides rich information that can be used to separate jump risks from volatility risks, but the identification of these distinct features hinges crucially on a suitable modeling framework that enables the jumps to exert a unique and distinct impact on the model-implied surface relative to the volatility factors.

Third, the Supplementary Appendix reports results for subsample estimation covering the periods 1996-2006 and 2007-2010. We find the fit to be excellent for the long 1996-2006 subsample, with an overall RMSE of 1.05%, and with violations at the 1% level for three of the six option categories to be around 1%, while the remaining rates range from 8% to 18.5%, with the hardest category to fit being the long maturity OTM put options. While this sample excludes the financial crisis and thus poses less of a challenge to the option pricing model, the sample still covers some dramatic episodes around 1997, 1998 and the internet bubble, along with some extremely low volatility levels during 2004-2006. Obviously, the shorter and extremely turbulent 2007-2010 period is much harder to accommodate, and the RMSE is now 2.36%, with violations of the fit to the six option categories ranging from 16% to 61% and the short maturity OTM call options being especially hard to price appropriately. Nonetheless, the similarities in the implied jump processes across the two subsamples and the full sample results for model (3.4) reflects a remarkable stability in the inference. The distinct decay parameters for the left and right tails are quite close and the mean jump sizes are between -4.5% and -5% for the negative tail and around 2% for the positive tail in both subsamples. Moreover, both imply a half life for volatility shocks of about 5-6 weeks for V_1 and about one year for V_2, which aligns well with the implied quantities for the full sample. Thus, the critical dynamic features of the estimated model are remarkably stable over time.

We conclude that the novel three-factor model (3.4) provides a significant improvement over the traditional specification of the risk-neutral equity return dynamics driving the evolution of the option surface. Furthermore, the more critical features of the system seem to be quite robustly estimated across highly diverse market conditions. As such, the system should, comparatively, provide a sound basis for the exploration of the equity and volatility risk premiums which we turn to in the remainder of the paper.

4 Risk Premia Dynamics and Predictability

In this section we relate our findings based on the option panel to the underlying return data. In particular, we study the links between the state vector, or risk factors, extracted from the option panel with the volatility and jump risks inferred from the dynamics of the underlying stock market index. This sets the stage for a direct exploration of the associated equity and variance risk premiums and their relation with the option-implied factors. Given the superior performance, we rely on the three factor model (3.4) for tracking the dynamics of the option surface over time. We continue to avoid making strong assumptions regarding the dynamics of the market risks under the actual probability measure. Consequently, this part of our analysis is fully nonparametric, and we invoke only minimal stationarity conditions regarding the $\mathbb{P}$-dynamics.

4.1 Connecting the Information in the Option Panel and the Underlying Asset

To define risk premia, we must first develop some consistent notation concerning the pricing of each of the sources of risks in the model (3.4). We define

$$\begin{aligned} W_{1,t}^{\mathbb{P}} &= W_{1,t}^{\mathbb{Q}} - \int_0^t \lambda_s^{W_1}\, ds, \quad W_{2,t}^{\mathbb{P}} = W_{2,t}^{\mathbb{Q}} - \int_0^t \lambda_s^{W_2}\, ds, \\ B_{1,t}^{\mathbb{P}} &= B_{1,t}^{\mathbb{Q}} - \int_0^t \lambda_s^{B_1}\, ds, \quad B_{2,t}^{\mathbb{P}} = B_{2,t}^{\mathbb{Q}} - \int_0^t \lambda_s^{B_2}\, ds, \end{aligned} \tag{4.1}$$

where $W_{1,t}^{\mathbb{P}}$, $W_{2,t}^{\mathbb{P}}$, $B_{1,t}^{\mathbb{P}}$ and $B_{2,t}^{\mathbb{P}}$ are $\mathbb{P}$ Brownian motions and $\lambda_t^{W_1}$, $\lambda_t^{W_2}$, $\lambda_t^{B_1}$ and $\lambda_t^{B_2}$ denote the associated prices of risk. The compensator of the jump measure, μ, under the $\mathbb{P}$ measure is given by $dt \otimes \nu_t^{\mathbb{P}}(dx, dy)$, and the mapping $\nu_t^{\mathbb{P}}(dx, dy) \rightarrow \nu_t^{\mathbb{Q}}(dx, dy)$, defined for every jump size x and y and every point in time t, reflects the compensation for jump risk.

The dynamics of the stock price process under the physical probability measure $\mathbb{P}$ is then,

$$\frac{dX_t}{X_{t-}} = \alpha_t\, dt + \sqrt{V_{1,t}}\, dW_{1,t}^{\mathbb{P}} + \sqrt{V_{2,t}}\, dW_{2,t}^{\mathbb{P}} + \int_{\mathbb{R}^2} (e^x - 1)\, \widetilde{\mu}^{\mathbb{P}}(dt, dx, dy), \tag{4.2}$$

where

$$\alpha_t - (r_t - \delta_t) = \lambda_t^{W_1}\sqrt{V_{1,t}} + \lambda_t^{W_2}\sqrt{V_{2,t}} + \int_{\mathbb{R}^2} (e^x - 1)\, \nu_t^{\mathbb{P}}(dx, dy) - \int_{\mathbb{R}^2} (e^x - 1)\, \nu_t^{\mathbb{Q}}(dx, dy), \tag{4.3}$$

is the instantaneous equity risk premium, reflecting compensation for diffusive and jump price risks.

Application of Itô formula yields the following representation for $\log(X_t)$ under $\mathbb{P}$,

$$d\log(X_t) = \left[\alpha_t - q_t^{\mathbb{P}}\right] dt + \sqrt{V_{1,t}}\, dW_{1,t}^{\mathbb{P}} + \sqrt{V_{2,t}}\, dW_{2,t}^{\mathbb{P}} + \int_{\mathbb{R}^2} x\, \widetilde{\mu}^{\mathbb{P}}(dt, dx, dy), \tag{4.4}$$

where $q_t^{\mathbb{P}} = \frac{1}{2}V_t + \int_{\mathbb{R}^2}(e^x - 1 - x)\,\nu_t^{\mathbb{P}}(dx,dy)$, and similarly under $\mathbb{Q}$ with α_t replaced by $r_t - \delta_t$ and all superscripts $\mathbb{P}$ replaced with $\mathbb{Q}$ in the expression above.

We may then define the (conditional) cum-dividend equity risk premium over the horizon τ,

$$\begin{aligned}\text{ERP}_t^\tau &\equiv \frac{1}{\tau}\left[\mathbb{E}_t^{\mathbb{P}}\left(\log\left(\frac{X_{t+\tau}}{X_t}\right) + \int_t^{t+\tau}(\delta_s + q_s^{\mathbb{P}})ds\right) - \mathbb{E}_t^{\mathbb{Q}}\left(\log\left(\frac{X_{t+\tau}}{X_t}\right) + \int_t^{t+\tau}(\delta_s + q_s^{\mathbb{Q}})ds\right)\right] \\ &\quad - \frac{1}{\tau}\left[\mathbb{E}_t^{\mathbb{P}}\left(\int_t^{t+\tau} r_s ds\right) - \mathbb{E}_t^{\mathbb{Q}}\left(\int_t^{t+\tau} r_s ds\right)\right] \\ &= \frac{1}{\tau}\,\mathbb{E}_t^{\mathbb{P}}\left(\int_t^{t+\tau}(\alpha_s - (r_s - \delta_s))\,ds\right).\end{aligned} \tag{4.5}$$

Notice that since a long position in the market index involves a commitment of capital a part of the wedge in the $\mathbb{P}$ and $\mathbb{Q}$ expectations of the cum-dividend equity returns reflects compensation for the time-variation in the risk-free interest rate.[12]

We next define corresponding risk measures for the return variation. The most natural and popular one is the quadratic variation over $[t, t+\tau]$ which we denote $QV_{t,t+\tau}$. This captures the return variation over the given horizon and is given by,

$$\begin{aligned}QV_{t,t+\tau} &= QV^c_{t,t+\tau} + QV^j_{t,t+\tau}, \\ QV^c_{t,t+\tau} &= \int_t^{t+\tau}(V_{1,s} + V_{2,s})\,ds, \\ QV^j_{t,t+\tau} &= \int_t^{t+\tau}\int_{\mathbb{R}^2} x^2\,\mu(ds,dx,dy),\end{aligned} \tag{4.6}$$

where we decompose the return variation into terms generated by the continuous and jump component of X, $QV^c_{t,t+\tau}$ and $QV^j_{t,t+\tau}$. Further, note that the (realized) quadratic variation is independent of the probability measure. The variance risk premium is defined as,

$$\text{VRP}_t^\tau \equiv \frac{1}{\tau}\left[\mathbb{E}_t^{\mathbb{P}}\left(QV_{t,t+\tau}\right) - \mathbb{E}_t^{\mathbb{Q}}\left(QV_{t,t+\tau}\right)\right], \tag{4.7}$$

which provides compensation for the variance risk associated with both the continuous and the jump component of X.

We are also interested in assessing directly the risks and risk premiums associated with jumps. In particular, we want to gauge the compensation for large price jumps and to allow for a separate risk premium for the negative versus positive jumps. We obtain direct measures of the jump risks by simply counting the number of "big" jumps over the relevant horizon,

$$LT^K_{t,t+\tau} \equiv \int_t^{t+\tau}\int_{\mathbb{R}^2} 1_{\{x\le -K\}}\,\mu(ds,dx,dy), \qquad RT^K_{t,t+\tau} \equiv \int_t^{t+\tau}\int_{\mathbb{R}^2} 1_{\{x\ge K\}}\,\mu(ds,dx,dy), \tag{4.8}$$

[12]This term is, of course, absent if we instead define the equity risk premium using futures on the market index.

where K is a prespecified threshold. We set $K = 0.5\%$ in the subsequent analysis.[13] From the properties of the compensator for a jump measure, we have,

$$\begin{aligned} LT^K_{t,t+\tau} &= \int_t^{t+\tau}\int_{\mathbb{R}^2} 1_{\{x\le -K\}}\,\nu^{\mathbb{P}}_s(dx,dy)\,ds + \epsilon^L_{t,t+\tau}, \quad \mathbb{E}^{\mathbb{P}}_t(\epsilon^L_{t,t+\tau}) = 0, \\ RT^K_{t,t+\tau} &= \int_t^{t+\tau}\int_{\mathbb{R}^2} 1_{\{x\ge K\}}\,\nu^{\mathbb{P}}_s(dx,dy)\,ds + \epsilon^R_{t,t+\tau}, \quad \mathbb{E}^{\mathbb{P}}_t(\epsilon^R_{t,t+\tau}) = 0. \end{aligned} \tag{4.9}$$

Hence, up to martingale difference sequences, $LT^K_{t,t+\tau}$ and $RT^K_{t,t+\tau}$ measure the $\mathbb{P}$ jump intensity of "large" jumps.

The risk measures we have introduced, including, $QV^c_{t,t+\tau}$, $QV^j_{t,t+\tau}$, $\widehat{LT}^K_{t,t+\tau}$ and $\widehat{RT}^K_{t,t+\tau}$, are not directly observable. However, we can estimate them using high-frequency futures data. We introduce the following empirical measures,

$$\begin{cases} \widehat{LT}^K_{t,t+\tau} = \displaystyle\sum_{i:\frac{i}{n}\in(t,t+\tau]} 1_{\{\Delta^n_i f > (K\vee\alpha n^{-\varpi})\}}, \\ \widehat{RT}^K_{t,t+\tau} = \displaystyle\sum_{i:\frac{i}{n}\in(t,t+\tau]} 1_{\{\Delta^n_i f < -(K\vee\alpha n^{-\varpi})\}}, \end{cases} \tag{4.10}$$

$$\widehat{QV}_{t,t+\tau} = \sum_{i:\frac{i}{n}\in(t,t+\tau]} |\Delta^n_i f|^2, \quad \widehat{QV}^c_{t,t+\tau} = \sum_{i:\frac{i}{n}\in(t,t+\tau]} |\Delta^n_i f|^2 1_{\{|\Delta^n_i f|\le \alpha n^{-\varpi}\}}, \tag{4.11}$$

where the tuning parameters α and ϖ are determined in the same way as for the estimation of $\widehat{V}^{(n,m_n)}_t$ in Section 2. Under fairly general conditions we can show, see, e.g., Jacod (2008),

$$\widehat{LT}^K_{t,t+\tau} \xrightarrow{\mathbb{P}} LT^K_{t,t+\tau}, \ \widehat{RT}^K_{t,t+\tau} \xrightarrow{\mathbb{P}} RT^K_{t,t+\tau}, \ \widehat{QV}_{t,t+\tau} \xrightarrow{\mathbb{P}} \widehat{QV}_{t,t+\tau}, \ \widehat{QV}^c_{t,t+\tau} \xrightarrow{\mathbb{P}} \widehat{QV}^c_{t,t+\tau}. \tag{4.12}$$

Similarly, we estimate the equity and variance risk premiums using the relations,

$$\begin{aligned} &\log\left(\frac{X_{t+\tau}}{X_t}\right) - \frac{1}{\tau}\int_t^{t+\tau}(r_s-\delta_s)\,ds = \mathrm{ERP}^\tau_t + \frac{1}{\tau}\mathbb{E}^{\mathbb{P}}_t\left(\int_t^{t+\tau} q^{\mathbb{P}}_s ds\right) + \epsilon^E_{t,t+\tau}, \quad \mathbb{E}^{\mathbb{P}}_t(\epsilon^E_{t,t+\tau}) = 0, \\ &\widehat{\mathrm{VRP}}^\tau_t = \frac{1}{\tau}\left[\widehat{QV}_{t,t+\tau} - \mathbb{E}^{\mathbb{Q}}_t\left(QV_{t,t+\tau}\right)\right] = \mathrm{VRP}^\tau_t + \epsilon^V_{t,t+\tau}, \quad \mathbb{E}^{\mathbb{P}}_t(\epsilon^V_{t,t+\tau}) = 0, \end{aligned} \tag{4.13}$$

where $\mathbb{E}^{\mathbb{Q}}_t\left(QV_{t,t+\tau}\right)$ can be measured in model-free fashion via the VIX index computed by the CBOE from a portfolio of S&P 500 index options. Equation (4.13) shows that a martingale difference sequence separates $\widehat{\mathrm{VRP}}^\tau_t$ from VRP^τ_t, and, likewise, a martingale difference sequence separates the log excess cum-dividend returns on the underlying asset from the unobservable $\mathrm{ERP}^\tau_t + \frac{1}{\tau}\mathbb{E}^{\mathbb{P}}_t\left(\int_t^{t+\tau} q^{\mathbb{P}}_s ds\right)$. In principle, we can remove the term stemming from the convexity

[13] This (fixed) threshold K is large enough that we can separate returns exceeding this (absolute) level from diffusive volatility using 1-minute observations. Experiments with alternative cutoffs produced similar results.

adjustment, i.e., $\frac{1}{\tau}\mathbb{E}_t^{\mathbb{P}}\left(\int_t^{t+\tau} q_s^{\mathbb{P}} ds\right)$, via a consistent estimator for $\int_t^{t+\tau} q_s^{\mathbb{P}} ds$ (again up to a martingale difference term) obtained from high-frequency data, e.g.,[14]

$$\int_t^{t+\tau} \widehat{q}_s^{\mathbb{P}}\, ds \;=\; \sum_{i:\frac{i}{n}\in(t,t+\tau]} \left[\frac{1}{2}\left|\Delta_i^n f\right|^2 1_{\{|\Delta_i^n f|\leq \alpha n^{-\varpi}\}} \;+\; \left(e^{\Delta_i^n f} - 1 - \Delta_i^n f\right) 1_{\{|\Delta_i^n f|>\alpha n^{-\varpi}\}}\right].$$

In practice, this adjustment is minute and the results are virtually unchanged if we implement it.[15]

4.2 The Predictability of Equity and Variance Risk and Risk Premia

Equipped with empirically feasible estimators for the relevant risk measures and risk premia, we may now explore the relationship between the option-implied factors, $V_{1,t}$, $V_{2,t}$ and U_t, that drive the dynamics of the option surface and the various risk measures and risk premia associated with the underlying asset. We rely on alternative versions of the following predictive regression,

$$y_t = \alpha_0 + \alpha_1 V_{1,t} + \alpha_2 V_{2,t} + \alpha_3 U_t + \epsilon_t, \tag{4.14}$$

where the left hand side variable represents, in turn, jump and diffusive variance risk measures as well as risk premia, i.e., $y_t = \widehat{LT}^K_{t,t+\tau}$, $\widehat{RT}^K_{t,t+\tau}$, $\widehat{QV}^c_{t,t+\tau}$, $\log\left(\frac{X_{t+\tau}}{X_t}\right) - \frac{1}{\tau}\int_t^{t+\tau}(r_s - \delta_s)ds$ and $\widehat{\text{VRP}}_t^{\tau}$. Given the relationships explicated in equations (4.10) and (4.13), it is evident that the regressions based on the alternative y_t variables, asymptotically, will yield the identical estimates to those based on the corresponding infeasible measures of interest, i.e., $\int_t^{t+\tau}\int_{\mathbb{R}} 1_{\{x\leq -K\}}\nu_s^{\mathbb{P}}(dx,dy)$, $\int_t^{t+\tau}\int_{\mathbb{R}} 1_{\{x\geq K\}}\nu_s^{\mathbb{P}}(dx,dy)$, $QV^c_{t,t+\tau}$, ERP_t^{τ} and VRP_t^{τ}, respectively. Thus, the predictive regression in equation (4.14) speaks directly to the linkages between the option surface dynamics and the risks and risk premia associated with the equity-index market.

In general, if the premia for the diffusive and jump risks are spanned by the three factors V_1, V_2 and U, then the expectation of y_t conditional on the time t information set, $\mathcal{F}_t$, in equation (4.14) should be functionally related to $V_{1,t}$, $V_{2,t}$ and U_t. Moreover, in the standard case, almost universally adopted in empirical option pricing applications, the measure change preserves the affine structure, so the conditional mean of y_t in equation (4.14) is a linear function of $V_{1,t}$, $V_{2,t}$ and U_t. Hence, the regression in (4.14) produces optimal (mean-square error) predictors for the volatility and jump risks using the information at time t. We further note that, conceptually, our extraction of the factors $V_{1,t}$, $V_{2,t}$ and U_t provides a richer information set for forecasting the relevant volatility and jump realizations than is available from the history of the underlying asset returns. The latter,

[14] For the overnight periods we just take one half of the squared return.

[15] We do not report results adjusting for this term to conserve space. They are available upon request.

at best, generates estimates of the path for $\{V_{1,s}+V_{2,s}\}_{s\leq t}$ and an associated measure of the return variation generated by jumps.[16]

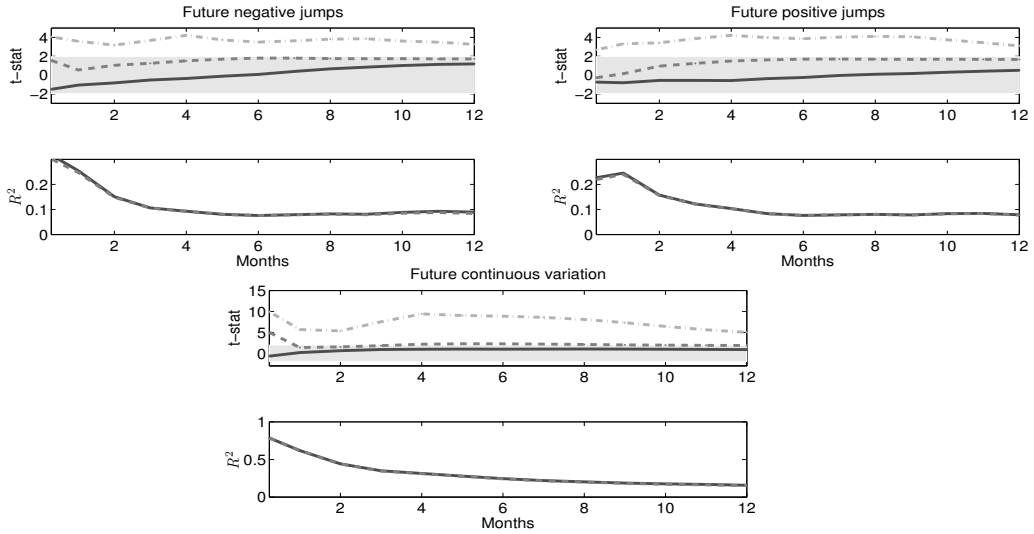

Figure 9: **Predictive regressions for volatility and jump risks.** The volatility and jump risk measures are defined in (4.10)-(4.11). For each regression, the top panels depict the t-statistics for the individual parameter estimates while the bottom panels indicate the regression R^2. The predictive variables are V_1 (dashed-dotted line), V_2 (dashed line) and $\widetilde{U}$ (solid line), where $\widetilde{U}$ is the residual from the linear projection of U on V_1 and V_2. The regression standard errors are constructed to also account for the estimation error in the projection generating $\widetilde{U}$. The dashed lines in the R^2 plots correspond to constrained regressions, including only V_1 and V_2.

The results from the predictive regressions are summarized on Figures 9 and 10. Given our relatively short sample, we compute the predictive regressions for horizons up to one year only. Our findings may be summarized as follows. First, Figure 9 shows that the state variables, extracted from the option panel, help predict the future evolution of risks. In particular, the plot pertaining to the count of positive and negative jumps demonstrates, convincingly, that the jump intensity varies under the statistical measure, $\mathbb{P}$, i.e., $\nu_t^{\mathbb{P}}(dx,dy)$ is truly a function of t.[17] At the same time, the three state variables V_1, V_2 and U have very different predictive power for the future volatility and jump risks of the underlying asset. Clearly, most of the predictability in Figure 9 stems from the volatility factors, V_1 and V_2. Once they are included in the regressions, U provides no incremental explanatory power for the future volatility and jump risk. This is evident both

[16]However, it is, of course, essential that we initially "correct" the option data for the relevant risk premia through the estimation of the parametric model (3.4).

[17]For further nonparametric evidence of time variation in the $\mathbb{P}$-jump intensity, using only high-frequency futures data, see Bollerslev and Todorov (2011a).

from the insignificant t-statistics corresponding to U and the negligible drop in the R^2 value when excluding U from the regressions.[18] Hence, Figure 9 is consistent with a model for which the jump intensity, under $\mathbb{P}$, depends only on the two volatility factors V_1 and V_2, but not U.[19]

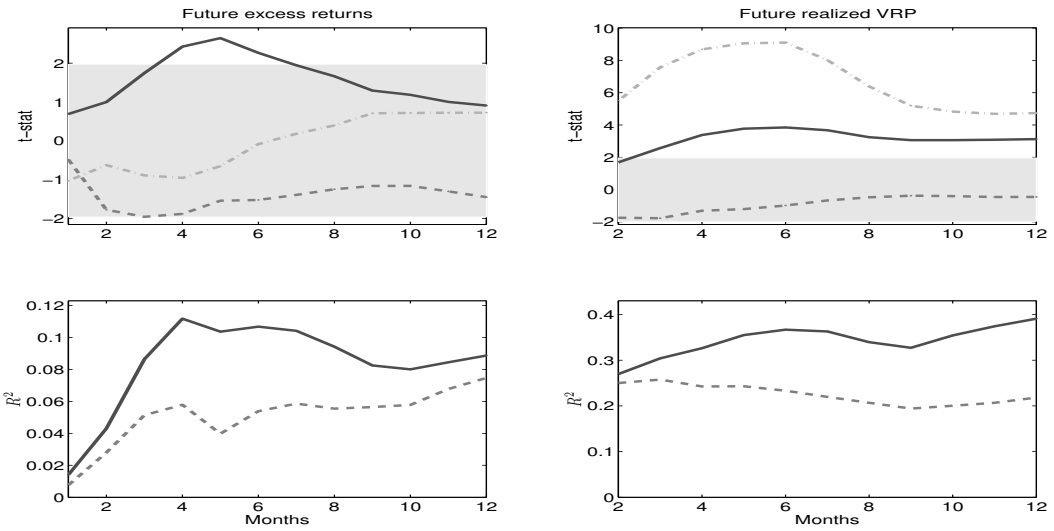

Figure 10: **Predictive regressions for equity and variance risk premia.** For each regression, the top panels depict the t-statistics for the individual parameter estimates while the bottom panels indicate the regression R^2. The predictive variables are V_1 (dashed-dotted line), V_2 (dashed line) and $\widetilde{U}$ (solid line), where $\widetilde{U}$ is the residual from the linear projection of U on V_1 and V_2. The regression standard errors are constructed to also account for the estimation error in the projection generating $\widetilde{U}$. The dashed lines in the R^2 plots correspond to constrained regressions, including only V_1 and V_2.

Turning to the predictability of the equity and variance risk premia, displayed on Figure 10, we see that the three factors, extracted from the option panel, now take on very different roles. In particular, a significant part of the predictability of both the equity and the variance risk premia is due to the third factor U. For the variance risk premia, V_1 also contributes strongly to the predictive power at shorter horizons, while, importantly, for the equity risk premium U is the dominant explanatory factor for all horizons up through at least nine months. The similarity between the predictive regression results for the equity and variance risk premia in terms of the role of U is consistent with Bollerslev and Todorov (2011b) who find that the equity and variance risk premia have a common component stemming from the compensation for left tail jump risk.

[18]This is consistent with Andersen and Bondarenko (2011) who find that moderate-to-deep OTM put options have no predictive power for the future volatility of the underlying equity index returns once they control for the information provided by ATM and OTM call options.

[19]In general, since V_1 contains jumps with a time-varying intensity that loads on all the factors, V_1, V_2 and U, the conditional forecast of future volatility $V_{1,t} + V_{2,t}$ still depends on the current value of the third factor U.

Recall that U appears only in the risk-neutral intensity of the negative jumps in model (3.4). The common dependence of the equity and variance risk premia on U, coupled with the significant persistence of the latter, help rationalize the predictive power of the variance risk premium for future excess returns, documented in Bollerslev et al. (2009) and Drechsler and Yaron (2011). Our finding of a limited role for the two volatility factors in this context is consistent with earlier work on the risk-return tradeoff, see, e.g., French et al. (1987) and Glosten et al. (1993).[20] Finally, we note that our results are not driven solely by the events surrounding the financial crisis. We obtain qualitatively identical results from subsamples that end in July 2007.[21]

To further highlight the importance of correct model specification when extracting information about risk premia from options, on Figures 11 and 12, we display the results from the same predictive regressions for the future risks and risk premia, but using the two-factor volatility model in (3.3). This model incorporates all main features included in traditional empirical option pricing models, with two volatility state variables driving the dynamics of the option surface, while allowing for both diffusive and jump leverage effects, co-jumps in returns and volatility, and separate decay rates for the left and right jump tails. Figure 11 shows that the results concerning the future evolution of volatility and jump risks are similar to the corresponding results on Figure 9 for our extended three-factor model (3.4). This is quite intuitive, given that the U factor plays only a minimal role in forecasting these quantities within model (3.4). Specifically, both volatility factors of model (3.3) help predict the occurrence of future (positive and negative) jumps as well as the future diffusive volatility. Likewise, the t-statistics and R^2 measures in Figures 9 and 11 are very similar. Turning to the predictive regressions for the risk premia, we find critical differences between the displays in Figures 10 and 12. First, for the equity risk premium, both volatility factors in model (3.3) are insignificant for most horizons, and when they appear marginally significant they enter with a *negative* sign. In addition, the R^2 of the predictive regression for the future excess equity returns is dramatically reduced in the two-factor model (3.3) relative to model (3.4). In short, the evidence for predictability of the equity risk premium vanishes when the option surface dynamics is modeled within the common jump-diffusive framework and driven exclusively by stochastic volatility factors. Turning to the predictive regressions for the variance risk premium, we find both volatility factors in model (3.3) to be significant. This differs from Figure 10 where, upon controlling for V_1 and U, the second volatility factor has no auxiliary predictive power for the future variance risk premium. Nonetheless, the R^2 of the variance risk premium regression is notably higher when we utilize the state variables from three-factor model (3.4) rather than those of the two-factor model (3.3).

[20] Predictive regressions for future excess returns involving only V_1 and V_2 produce insignificant t-statistics for both.
[21] The documentation of these results are available upon request.

Hence, the inclusion of the U factor, allowing for the left risk-neutral tail to have a separate source of variation, enables us to capture predictability in the equity risk premium and a substantial part of the variance risk premium that cannot be accounted for through – even an elaborate version of – the standard modeling framework. These findings have important implications for the assessment of how well structural economic models can rationalize the predictability of the equity and variance risk premia which we discuss next.

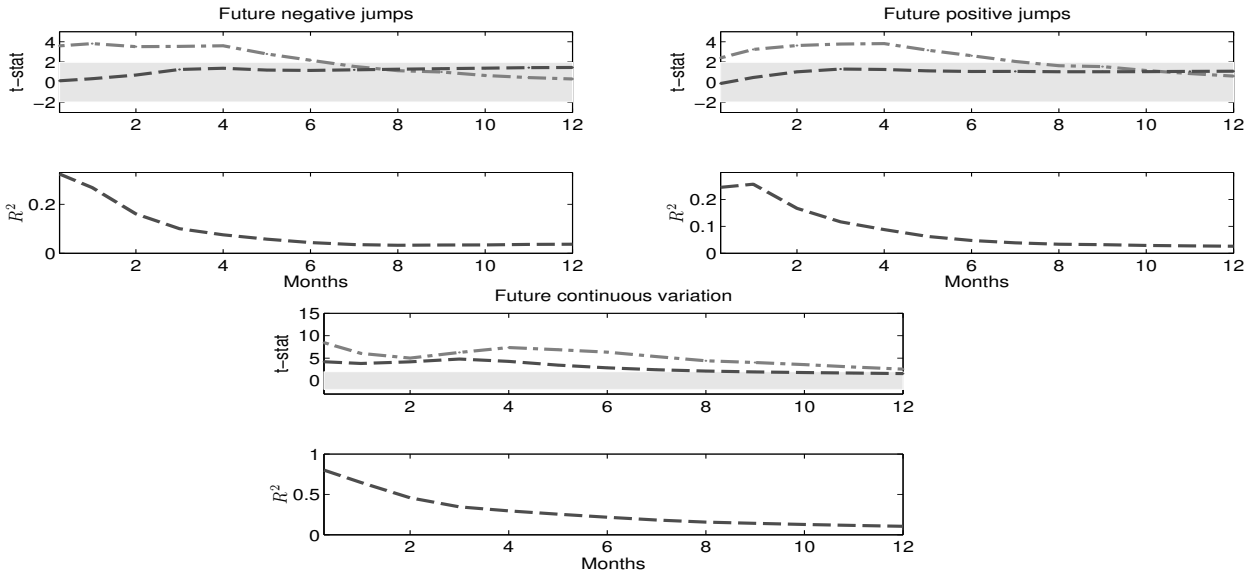

Figure 11: **Predictive regressions for volatility and jump risks in the two-factor model (3.3).** The predictive variables are V_1 (dashed-dotted line) and V_2 (dashed line). The remainder of the notation follows the conventions in Figure 9.

4.3 Structural Implications of the Predictive Power of the Option Surface

Figures 9 and 10 demonstrate that the factor U, driving a substantial part of the OTM short maturity put option dynamics, has no impact on the *actual* volatility and jump dynamics of the underlying asset. In contrast, the factor has a critical effect on the *pricing* of volatility and jump risk. In other words, it resembles a risk premium and not a risk factor. Can we rationalize this finding from an economic perspective? To guide intuition, we compare the findings in Figures 9 and 10 with those from recent structural models that link option prices to fundamental macroeconomic risks, such as aggregate consumption and dividends. This emerging literature has made important progress in tackling the challenging, yet critical, task of jointly explaining the equity return and risk premium dynamics in a coherent general equilibrium setting. For concreteness, we initially explore

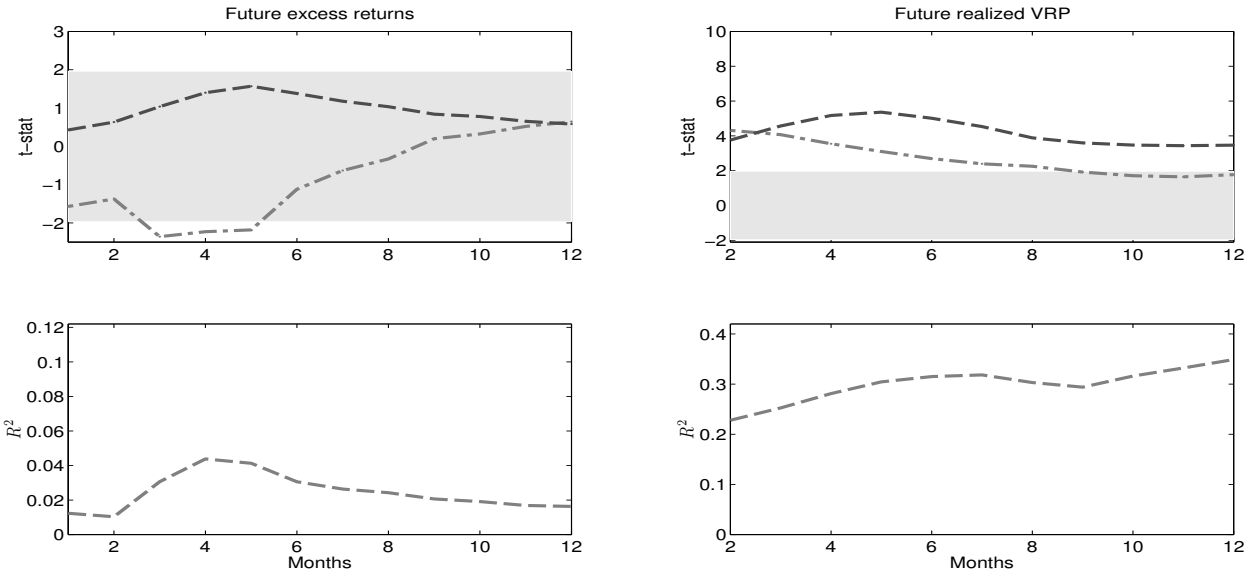

Figure 12: **Predictive regressions for the risk premia in the two-factor model (3.3)**. The predictive variables are V_1 (dashed-dotted line) and V_2 (dashed line). The remainder of the notation follows the conventions in Figure 10.

a couple of specific models featuring a representative agent with Epstein-Zin preferences exposed to risks in real consumption growth, namely Wachter (2013) and Drechsler and Yaron (2011).

In the one-factor model of Wachter (2013), the consumption growth is subject to infrequent, but large, negative jumps (rare disasters) with a time-varying arrival rate, resembling the mechanism in Gabaix (2012).[22] This type of equilibrium model can account for many critical empirical features such as the correlation between volatility and jump risks, the time-varying jump arrival as well as the ability of the market variance risk premium to predict future equity excess returns.

The model of Drechsler and Yaron (2011) specifies consumption growth as conditionally Gaussian with a time-varying conditional mean (long-run risk) and conditional volatility.[23] The system has three factors, with one driving the conditional mean and two governing the conditional volatility of consumption growth.[24] Drechsler and Yaron (2011) show that the time-varying jump intensity – in turn governed by the volatility state – explains a significant part of the predictability in the variance risk premia of future excess returns.

Figures 13 and 14 summarize the findings from predictive regressions for future volatility and

[22]These papers generalize work by Barro (2006) and Barro and Ursua (2008) in which rare disasters are i.i.d.

[23]This model builds on Eraker and Shaliastovich (2008) and generalizes many prior models in which consumption growth contains a small predictable persistent component, including the original work of Bansal and Yaron (2004).

[24]One of the state variables that we label volatility factors directly controls the conditional variance of consumption growth, while the other captures the variation in the long run variance of consumption growth.

jump risks as well as equity and variance risk premia in the models of Wachter (2013) and Drechsler and Yaron (2011), using the concurrent level of the relevant state variables in each model as predictors. Not surprisingly, given the one-factor structure, the Wachter (2013) model faces some challenges in accommodating the evidence laid out in Figures 9 and 10. Importantly for our analysis, the predictability of future excess returns and variance risk premia is linked closely to the predictability of the future return variation, and the underlying pattern of significance is identical in all cases (flat line), and the degree of explanatory power rises roughly linearly with maturity. Compared to Figure 9, the predictability in the Wachter (2013) model is inverted, as the return variation is forecast with relatively higher precision over long rather than short horizons. Furthermore, the explanatory power is uniformly too low. Likewise, referencing Figure 10, the model fails to capture the degree and pattern of predictability in the excess returns and variance risk premium.[25]

Turning to Figure 14, it is evident that the more flexible volatility structure of Drechsler and Yaron (2011) is useful in accommodating some of the stylized features of the data. Nonetheless, it is equally clear that the long-run risk factor helps predict neither the future volatility and jump risks nor the equity and variance risk premia. In this structural setting, essentially all predictability stems from the two volatility factors. They provide the channel through which past variance risk premia generate predictable movements in the equity risk premium. Thus, relative to our empirical finding, captured by Figures 9 and 10, this structural model also ties the predictability of future volatility and jump risks too closely to the predictability of the equity and variance risk premia. Equivalently, the structural model implies a tight relationship between the dynamics of the option panel and the return dynamics of the underlying equity market. In contrast, our empirical results based on model (3.4) – nesting all standard two-factor affine volatility specifications – document a partial, and critical, decoupling between the factors driving the equity return dynamics and those governing the pricing of risk, and thus the equity and variance risk premia.

There is a fundamental reason for the discrepancy between our empirical findings and the implications of the structural models of Wachter (2013), and its extension in Seo and Wachter (2013), as well as Drechsler and Yaron (2011). Although the models generate risk premia through different channels – the presence of rare disasters and uncertainty about their arrival (Wachter (2013) and Seo and Wachter (2013)) versus long-run risk and stochastic volatility in consumption growth (Drechsler and Yaron (2011)) – they share a critical feature in the pricing of jump tail

[25]A two-factor extension, like in Seo and Wachter (2013), in which rare disaster probability is driven by two factors, can potentially generate dynamic patterns more consistent with the data. However, as we discuss below in more general terms, such an extension will still produce a close link between the predictability of future risks and risk premia unlike what we find in the data.

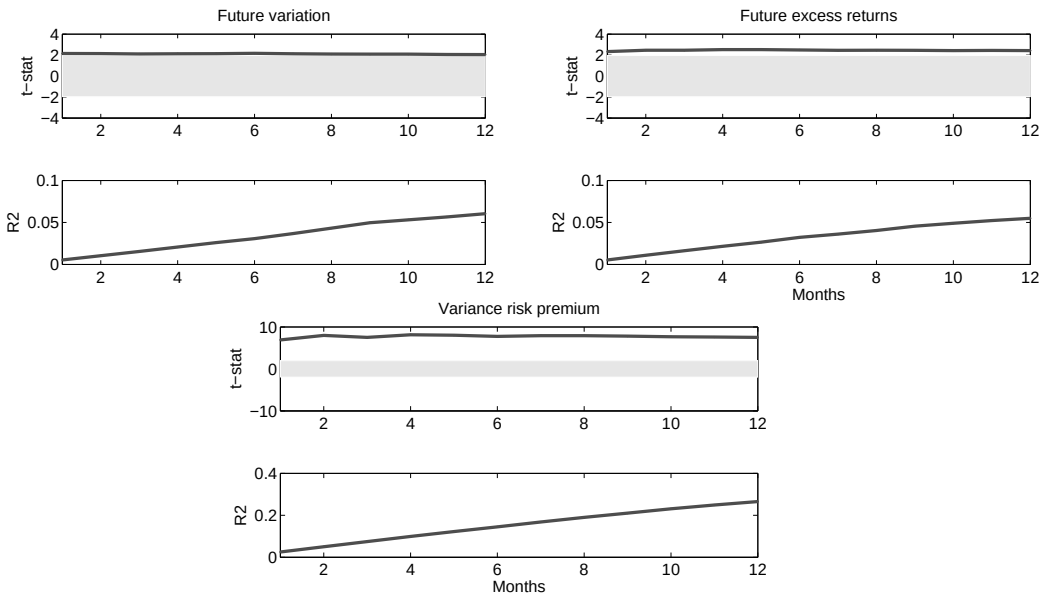

Figure 13: **Predictive regressions implied by the Wachter (2013) structural model.** The predictive variable is time-varying intensity of rare disaster in consumption growth.

risk. They both imply that the ratio $\frac{\nu_t^{\mathbb{Q}}(dx,dy)}{\nu_t^{\mathbb{P}}(dx,dy)}$ is time-invariant. That is, the risk-neutral jump intensity is proportional to that under the actual probability measure, so the two jump intensities are equivalent in terms of their time variation. Therefore, the jump risk premia are generated by changing the distribution of the jump size only. This implies that the variation in the jump intensity is "inherited" under the equivalent change of measure. Consequently, these equilibrium based pricing kernels cannot "generate" new state variables in addition to those that drive the fundamental risks in the economy. In turn, this necessarily generates the tight link between the dynamics of the underlying asset and the option surface within these models.

In fact, this tight linkage of the physical and risk-neutral jump intensities is operative for a wide class of popular structural models with representative agents having Epstein-Zin preferences. The part of the density $\frac{d\mathbb{Q}}{d\mathbb{P}}$ due to the change of the jump measure is characterized by,

$$\mathcal{E}\left(\int_0^t \int_{\mathbb{R}^n} (Y(\omega,s,\mathbf{x})-1)\widetilde{\mu}^{\mathbb{P}}(ds,d\mathbf{x})\right), \tag{4.15}$$

where the jump component of the state vector in the economy under the $\mathbb{P}$ probability measure is given as a (multivariate) integral of the form $\int_0^t \int_{\mathbb{R}^n} \mathbf{x}\,\widetilde{\mu}^{\mathbb{P}}(ds,d\mathbf{x})$; $Y(\omega,t,\mathbf{x}) = \frac{\nu_t^{\mathbb{Q}}(d\mathbf{x})}{\nu_t^{\mathbb{P}}(d\mathbf{x})}$ is the measure change for the jump intensity, and $\mathcal{E}$ is the Doleans-Dade exponential.[26] In general, $Y(\omega,t,\mathbf{x})$

[26]For the expression in (4.15) and the definition of the Doleans-Dade exponential, see Jacod and Shiryaev (2003),

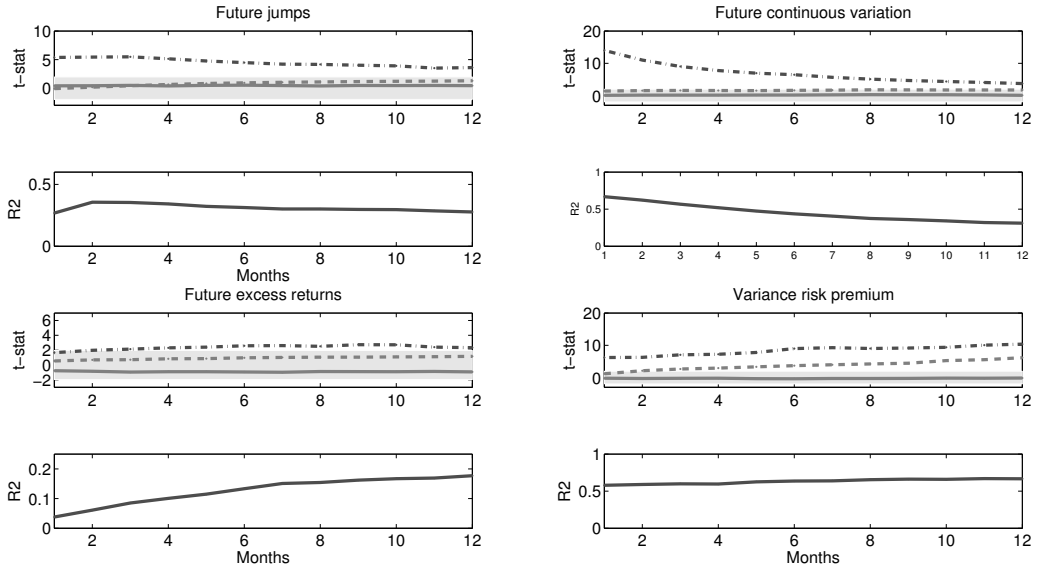

Figure 14: **Predictive regressions implied by the Drechsler and Yaron (2011) structural model.** The predictive variables are the conditional mean of consumption growth (solid line), stochastic volatility of consumption growth (dashed-dotted line) and central tendency of stochastic volatility (dashed line). The dashed lines in the R^2 plots correspond to constrained regressions including only the volatility state variables.

will be stochastic. However, within an equilibrium setting, stipulating an affine dynamics for the fundamentals along with a representative agent with Epstein-Zin preferences generates the restriction that the expression in equation (4.15) is exponentially affine in the state vector, see, e.g., equation (2.22) of Eraker and Shaliastovich (2008). In turn, this implies that $Y(\omega, s, \mathbf{x})$ must be non-random and time-invariant, i.e., depend only on $\mathbf{x}$.[27] Hence, the equilibrium implied statistical and risk-neutral intensities of the price and volatility jumps (which are mixtures of the jumps in the state variables driving the fundamentals in the equilibrium model) will be affine functions of the *same* state vector.[28] In contrast, for our extended three-factor model, in which $\nu_t^{\mathbb{P}}(\cdot)$ loads, at most, very marginally on U_{t-}, there is a natural wedge between the time variation in statistical and risk-neutral jump intensities.

There are several ways in which the link between the asset and option price dynamics may

Corollary III.5.22 and I.4.59, respectively.

[27]The restriction on $Y(\omega, s, \mathbf{x})$ is actually stronger. In the equilibrium model, $Y(\omega, s, \mathbf{x})$ is an exponential function of jump size, see Theorem 1 of Eraker and Shaliastovich (2008). Hence, the jump measure change is implemented by exponential tilting (in Laplace transform space), with the degree of tilting determined by the preference parameters of the representative agent.

[28]In the models of Wachter (2013), Seo and Wachter (2013) and Drechsler and Yaron (2011), the jump intensities are even more tightly connected, as they are directly proportional.

be relaxed within an equilibrium setting to potentially account for our empirical evidence. They all involve generalizing the preferences of the representative agent in some form. One approach is to allow the representative agent's coefficient of risk aversion to vary over time. Du (2010) proposes a generalization of a habit formation model in which consumption growth is i.i.d. with rare jumps, building on the equilibrium models of Campbell and Cochrane (1999) and Barro (2006). In this model, and consistent with our findings, there is a wedge between the time-variation of the $\mathbb{P}$ and $\mathbb{Q}$ jump intensities: the former is constant while the latter is time-varying.[29] However, the time variation of the $\mathbb{Q}$ jump intensity is driven by the habit formation mechanism and is purely a function of the consumption surplus ratio. The latter, however, determines uniquely the diffusive volatility as well. This is at odds with our findings as the presence of U in the risk-neutral jump intensity drives a wedge between the time variation of the latter and that of the stochastic volatility. Thus, it remains an open question whether a model with external habit formation can decouple the option and asset price dynamics in a manner reminiscent of our empirical findings. It is also unclear whether the frequency and intensity of stock market jumps can be mapped into corresponding jumps in the consumption growth rate as implied by this model of habit formation.

An alternative way to relax the link between the option and asset price dynamics is by recognizing that agents do not directly observe the state vector and therefore need to filter the states from observables. The absence of perfect information plus aversion to ambiguity (particularly about extreme negative risks), like in Hansen and Sargent (2008), or lack of confidence in the estimate of the state vector, like in Bansal and Shaliastovich (2010), can make investors appear more risk averse than they would be in a perfect information setting. The key question is whether the ambiguity aversion or confidence risk variation can generate the required gap between the dynamics of the statistical and risk-neutral jump intensities.[30] For example, the tight link between the jump intensity under $\mathbb{P}$ and $\mathbb{Q}$ remains within the generalization of the Drechsler and Yaron (2011) model, developed by Drechsler (2013), in which agents are ambiguous about part of the dynamics, including the jumps in the conditional mean and variance of consumption growth. The representative agent's ambiguity drives an additional wedge between fundamental risks and asset prices and helps explain why variance risk premia have superior predictive power, relative to volatility itself, for future returns. Nonetheless, the linearity of the pricing kernel with respect to the state vector implies that no new state variable is "generated" going from $\mathbb{P}$ to $\mathbb{Q}$, thus ultimately rendering the model predictions incompatible with our empirical findings.

[29] Our empirical findings, unlike the model of Du (2010), suggest strong time variation in the $\mathbb{P}$ jump intensity.

[30] Recall that, according to the estimates for (3.4), U does not covary perfectly with stochastic volatility.

5 Conclusion

We document that the standard exponentially-affine jump-diffusive specifications used in the empirical option pricing literature are incapable of fitting critical features of the option surface dynamics for the S&P 500 index, especially in scenarios involving significant shifts in the volatility smirk. We extend the risk-neutral volatility model to include a separate state variable which is crucial in capturing the time variation of priced downside tail risk. This new factor has no incremental explanatory power, beyond the traditional volatility factors, for the future evolution of volatility and jump risks. On the other hand, relative to the volatility components, the new factor provides critical, and superior, information for the time variation in the equity and variance risk premia.

Our findings demonstrate that the pricing in the option market is closely integrated with the underlying asset market. Moreover, the option panel embodies critical information regarding equity risk pricing that cannot be extracted directly from the underlying asset price dynamics. The wedge between the two probability measures arises primarily from the varying degree of compensation for downward tail jump risk. Our results suggest that time-varying risk aversion and/or ambiguity aversion, driven in part by the presence of large shocks, must be incorporated into structural asset pricing models if they are to explain the joint dynamics of the equity and option markets.

References

Andersen, T. G. and O. Bondarenko (2011). Dissecting the Market Pricing of Volatility. Working paper, Northwestern University and University of Illinois at Chicago.

Andersen, T. G., N. Fusari, and V. Todorov (2012). Parametric Inference and Dynamic State Recovery from Option Panels. Working paper, Northwstern University.

Bansal, R. and I. Shaliastovich (2010). Confidence Risk and Asset Prices. *American Economic Review 100*, 537–541.

Bansal, R. and A. Yaron (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance 59*, 1481–1509.

Barro, R. and J. F. Ursua (2008). Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity*, 255–335.

Barro, R. J. (2006). Rare Disasters and Asset Markets in the Twentieth Century. *Quarterly Journal of Economics 121*, 823–866.

Bates, D. S. (2000). Post-'87 Crash Fears in S&P 500 Future Options. *Journal of Econometrics 94*, 181–238.

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 (2011a). Estimation of Jump Tails. *Econometrica 79*, 1727–1783.

Bollerslev, T. and V. Todorov (2011b). Tails, Fears and Risk Premia. *Journal of Finance 66*, 2165–2211.

Bollerslev, T. and V. Todorov (2013). Time Varying Jump Tails. *Journal of Econometrics, forthcoming*.

Broadie, M., M. Chernov, and M. Johannes (2009). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance 62*, 1453–1490.

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. and L. Wu (2003). What Type of Process Underlies Options? A Simple Robust Test. *Journal of Finance 58*, 2581–2610.

Drechsler, I. (2013). Uncertainty, Time-Varying Fear, and Asset Prices. *Journal of Finance, forthcoming*.

Drechsler, I. and A. Yaron (2011). What's Vol Got to Do with It? *Review of Financial Studies 24*, 1–45.

Du, D. (2010). General Equilibrium Pricing of Options with Habit Formation and Event Risks. *Journal of Financial Economics 99*, 400–426.

Duffie, D., D. Filipović, and W. Schachermayer (2003). Affine Processes and Applications in Finance. *Annals of Applied Probability 13(3)*, 984–1053.

Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica 68*, 1343–1376.

Durrleman, V. (2008). Convergence of at-the-money Implied Volatilities to the Spot Volatility. *Journal of Applied Probability 45*, 542–550.

Eraker, B. and I. Shaliastovich (2008). An Equilibrium Guide to Designing Affine Pricing Models. *Mathematical Finance 18*, 519–543.

French, K., W. Schwert, and R. Stambaugh (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics 19*, 3–29.

Gabaix, X. (2012). Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance. *Quarterly Journal of Economics 127*, 645–700.

Glosten, L., R. Jaganathan, and D. Runkle (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance 48*, 1779–1801.

Hansen, L. and T. Sargent (2008). *Robustness*. Princeton University Press.

Jacod, J. (2008). Asymptotic Properties of Power Variations and Associated Functionals of Semimartingales. *Stochastic Processes and their Applications 118*, 517–559.

Jacod, J. and A. N. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.

Mancini, C. (2009). Non-parametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps. *Scandinavian Journal of Statistics 36*, 270–296.

Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics 63*, 3–50.

Seo, S. and J. A. Wachter (2013). Option Prices in a Model with Stochastic Disaster Risk. Working paper, University of Pennsylvania.

Singleton, K. (2006). *Empirical Dynamic Asset Pricing*. Princeton University Press.

Todorov, V. (2011). Econometric Analysis of Jump-Driven Stochastic Volatility Models. *Journal of Econometrics 160*, 12–21.

Wachter, J. A. (2013). Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *Journal of Finance 68*, 987–1035.