

The Pricing of Tail Risk and the Equity Premium: Evidence from International Option Markets

Torben G. Andersen, Nicola Fusari and Viktor Todorov

CREATES Research Paper 2018-2

The Pricing of Tail Risk and the Equity Premium: Evidence from International Option Markets*

Torben G. Andersen[†] Nicola Fusari[‡] Viktor Todorov[§]

November 2017

Abstract

We explore the pricing of tail risk as manifest in index options across international equity markets. The risk premium associated with negative tail events displays persistent shifts, unrelated to volatility. This tail risk premium is a potent predictor of future equity returns, while option-implied volatility only forecasts the future return variation. Hence, compensation for negative jump risk is the primary driver of the equity premium across all indices, whereas the reward for pure diffusive variance risk is largely unrelated to future equity returns. We also document pronounced commonalities, suggesting a high degree of integration among the major global equity markets.

Keywords: Equity Risk Premium, International Option Markets, Predictability, Tail Risk, Variance Risk Premium.

JEL classification: G12, G13, G15, G17.

*Andersen gratefully acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The work is partially supported by NSF Grant SES-1530748. We thank participants at the following presentations for useful comments: AQR Capital Management; Midwest Finance Association Meeting 2016, Chicago; Key-note address at the 2016 Nippon Finance Association Meeting, Yokohama; 30th Anniversary of GARCH Models and Measures Conference, Toulouse School of Economics; Key-note address at the Financial Econometrics and Empirical Asset Pricing Conference 2016, Lancaster University Management School; SoFiE Annual Meeting 2016, City University, Hong Kong; NBER Summer Institute 2016; Financial Econometrics Conference 2016, Questrom School of Business, Boston University; CBOE Conference on Volatility and Derivatives 2016, Chicago; Financial Econometrics Conference 2016, Rio de Janeiro; as well as participants at numerous seminars.

[†]Kellogg School, Northwestern University, Evanston, IL 60208, email: t-andersen@northwestern.edu.

[‡]The Johns Hopkins University Carey Business School, Baltimore, MD 21202, email: nfusari1@jhu.edu.

[§]Kellogg School, Northwestern University, Evanston, IL 60208, email: v-todorov@northwestern.edu.

1 Introduction

The last decade has witnessed a great deal of turmoil in global equity markets. These events represent a major challenge for dynamic asset pricing models. Can these models accommodate the observed interdependencies between volatility and tail events, and their pricing? How do volatility and tails differ across markets? How do the variance and tail risk premiums relate to the compensation for exposure to equity risk, i.e., the equity premium? The increasing liquidity of derivative markets worldwide provides an opportunity to shed light on these questions. In particular, the trading of equity-index options has grown sharply in the global financial centers, with both more strikes per maturity and additional maturities on offer. The latest development is a dramatic increase in the trading of options with short tenor. As a result, we now have access to active prices and quotes for financial securities that embed rich information about the pricing of market tail risk in many separate countries. In the current work, we draw on daily observations for option indices in the U.S. (S&P 500), Euro-zone (ESTOXX), Germany (DAX), Switzerland (SMI), U.K. (FTSE), Italy (MIB), and Spain (IBEX) over 2007-2014 to extract factors that are critical for the pricing of equity market risk across North America and Europe.

The confluence of turbulent periods renders recent years an excellent laboratory for analysis of the way investors treat evolving financial risks, and especially their attitude towards tail risk. Over our sample, three major shocks roiled the global markets, namely the global financial crisis originating in October 2008 and the European debt crises in May 2010 and August 2011. We exploit the option data to study how tail risk was perceived and priced across these episodes. By combining the pricing of financial risks, embodied within the equity-index option surfaces, with ex post information on realized returns, volatilities and jumps, we can gauge the relative size of the risk premiums and what factors drive the risk compensation. The performance of the individual indices varies drastically, with the German market appreciating by an average of about 8% per year and the Italian depreciating by 5% annually, as illustrated in Figure 1. We exploit this heterogeneity across countries to obtain additional power in studying the connection between market tail risk and the equity risk premium.

Our analysis is based on a parametric risk-neutral return model, but with no structural assumptions regarding the return data generating process. In this manner, we accommodate general risk premium dynamics which, in particular, are not tied closely to the level of market volatil-




Figure 1: **Country-specific cumulative equity-index cum dividend log-returns.** The source of the data used in constructing the figure are detailed in Section 3.

ity. The latter approach is traditional within common ICAPM inspired models, featuring a more direct risk-return trade-off. Specifically, standard option pricing models capture the dynamics of the equity-index option surface through the evolution of factors that determine the volatility of the underlying stock market, see, e.g., Bates (1996, 2000, 2003), Pan (2002), Eraker (2004) and Broadie et al. (2007).

However, recent evidence suggests that the fluctuations in the left tail of the risk-neutral density, extracted from equity-index options, cannot be spanned neither by regular volatility factors nor by realized risk measures constructed from return data. Hence, a distinct factor is necessary to account for the priced downside risk in the option surface, see Andersen et al. (2015b).¹

The European samples are shorter and contain fewer options in the strike cross-section, rendering separate day-by-day identification of the factors more challenging than for the U.S. Consequently, we specify a risk-neutral model with a single volatility component along with the tail factor. This facilitates robust factor identification and reduces the number of parameters, providing a solid basis for out-of-sample exploration of the predictive power of the option-implied factors.

Our empirical findings can be summarized as follows. In terms of return volatility, we find the European and U.S. markets to evolve in near unison through the financial crisis of 2008-2009. In contrast, we observe divergences in volatility during and after the initial European sovereign debt

¹Other models that have allowed for separate jump and volatility factors include Santa-Clara and Yan (2010), Christoffersen et al. (2012), Gruber et al. (2015), and Li and Zinna (2017).

crisis. Overall, the U.K., Swiss and, to some extent, German volatilities remain close throughout the sample. The largest divergences in volatility dynamics occur between the U.S., U.K. and Swiss indices on one side and the Spanish and Italian on the other, with the latter representing the Southern European countries in our sample.

For the left tail factor, there are interesting commonalities and telling differences. Again, the main divergences are associated with the Southern European indices, but even here there are, at times, stark discrepancies. Specifically, the Spanish tail factor reacts strongly to both phases of the sovereign debt crisis, while the Italian response is more muted, especially for the first phase.

We find a substantial gap between the time series evolution of priced tail risk and the level of market volatility for every option market we analyze. Moreover, a common feature emerges in the aftermath of crises: the left tail factor is correlated with volatility, yet it remains elevated long after volatility subsides to pre-crisis levels. This feature is incompatible with the usual approach to the modeling of volatility and jump risk in the literature.

The stark separation of the tail and volatility factors has important implications for the pricing and dynamics of market risk. Although volatility is a strong predictor of future market risk, such as the jump intensity and overall return variation, it provides no forecast power for the equity risk premium. In contrast, the component of the tail factor unspanned by volatility, the “pure tail factor,” has highly significant explanatory power for future returns. This factor also constitutes the primary driver of the negative tail risk premium, suggesting this is the operative channel through which it forecasts the equity premium. In particular, following crises, equity prices are heavily discounted and the option-implied tail factor remains elevated, even as market volatility resides. This combination provides a strong signal that the market will outperform in the future.

At the same time, there is evidence that the compensation for variance risk (the variance risk premium) is large (Carr and Wu (2009)) and forecasts future equity returns (Bollerslev et al. (2009)).² Given our finding that spot volatility predicts future return variation, but not future equity returns, we take a closer look at the source of explanatory power in the variance risk premium. We decompose the variance risk premium (VRP) into compensation for diffusive and jump risks. We document a high degree of correlation between the option-implied tail factor and the jump

²Bollerslev et al. (2009) studies the U.S. equity market only. Subsequent work has explored the predictive ability of the variance risk premium constructed from different underlying instruments as well as for predicting future excess returns on a variety of asset classes, e.g., credit spreads (Wang et al. (2013)), currency returns (Londono and Zhou (2017)), bond returns (Mueller et al. (2017)) and international stock market returns (Bollerslev et al. (2014)).

component of the VRP. On the contrary, the volatility factor is highly correlated with the diffusive component of the VRP, while it is uncorrelated with this jump factor. This suggests that the VRP provides explanatory power for the equity risk premium only due to the inclusion of the negative jump component. We confirm this hypothesis directly: the negative jump risk premium has better forecast power for future equity returns than the VRP, and the continuous component of the latter has negligible predictive power. That is, once the jump risk premium is stripped from the VRP, it no longer provides a signal regarding the future performance of the equity market. These findings are consistent across the sample period and all the equity indices.

We further document very strong dependence between the negative jump risk premiums for the equity indices. This relation is pronounced throughout the sample. By contrast, the dependence between the diffusive risk premiums across markets is very strong only for the financial crisis period 2007-2009, while for 2010-2014, associated with the European sovereign debt crises, the dependence between the diffusive risk premiums weakens significantly.

The rest of the paper is organized as follows. Section 2 presents the model we use to fit the option surfaces. We review the data used in the paper in Section 3. Section 4 reviews the estimation method, and demonstrates how we extract information and obtain the model-implied factors. Section 5 focuses on the option-implied factors. Section 5.1 displays the empirically extracted volatility and tail factors and discusses their interdependencies across indices. Section 5.2 explores the predictive power of the option-implied factors for future equity-index risk and returns. In Section 6, we focus on the interaction among the option-implied factors and the variance, left jump tail and equity risk premiums, with particular emphasis on the role of the pure jump factor. An analysis of the common features in the option-implied factors and risk premiums across indices is presented in Section 7. Finally, Section 8 concludes. Additional details on a variety of aspects related to the data, estimation method and results are contained in the Appendix.

2 Model

We start by introducing our model for the risk-neutral dynamics of a generic equity market index denoted by X . The model is designed to capture the salient dynamic features of the option surface through a low-dimensional state vector in a manner consistent with a continuous-time no-arbitrage setting. We will subsequently use the option-implied state vector to predict future equity risks

and risk premiums and relate the evidence to sources of priced risk in the economy. Our two-factor model for the risk-neutral index dynamics is given by the following restricted version of the representation in Andersen et al. (2015b),

$$\begin{aligned}\frac{dX_t}{X_{t-}} &= (r_t - \delta_t) dt + \sqrt{V_t} dW_t^{\mathbb{Q}} + \int_{\mathbb{R}} (e^x - 1) \tilde{\mu}^{\mathbb{Q}}(dt, dx), \\ dV_t &= \kappa_v (\bar{v} - V_t) dt + \sigma_v \sqrt{V_t} dB_t^{\mathbb{Q}} + \mu_v \int_{\mathbb{R}} x^2 1_{\{x < 0\}} \mu(dt, dx), \\ dU_t &= -\kappa_u U_t dt + \mu_u \int_{\mathbb{R}} x^2 1_{\{x < 0\}} \mu(dt, dx),\end{aligned}\tag{1}$$

where $(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}})$ is a two-dimensional Brownian motion with $\text{corr}(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}}) = \rho$, while μ denotes an integer-valued measure counting the jumps in the index, X , as well as the state vector, (V, U) , representing the spot variance and negative jump intensity. We denote the corresponding jump compensator by $dt \otimes \nu_t^{\mathbb{Q}}(dx)$, so the difference, $\tilde{\mu}^{\mathbb{Q}}(dt, dx) = \mu(dt, dx) - dt \nu_t^{\mathbb{Q}}(dx)$, constitutes the associated martingale jump measure. Finally, the drift term equals the risk-free interest rate minus the (continuous) dividend payout ratio, ensuring that the expected instantaneous return (under the risk-neutral measure) of the cum-dividend price equals the risk-free rate.

The jump component, x , captures price jumps, but also scenarios involving co-jumps. Specifically, for negative price jumps of size x , the two state variables, V and U , display (positive) jumps proportional to x^2 . Thus, the jumps in the spot variance and negative jump intensity are co-linear, albeit with distinct proportionality factors, μ_v and μ_u . This specification involves a substantial amplification from the negative price shocks to the risk factors. The compensator characterizes the conditional jump distribution and takes the form,

$$\frac{\nu_t^{\mathbb{Q}}(dx)}{dx} = U_t \cdot 1_{\{x < 0\}} \lambda_- e^{-\lambda_- |x|} + c_0^+ \cdot 1_{\{x > 0\}} \lambda_+ e^{-\lambda_+ x}.\tag{2}$$

The right hand side refers to negative and positive price jumps, respectively. Following Kou (2002), we assume that the price jumps are exponential, with separate tail decay parameters, λ_- and λ_+ , for negative and positive jumps. Finally, the left jump intensities is governed by the factor U_t while the positive jump intensity is constant and equal to c_0^+ .³

Our jump modeling includes novel features and deviates markedly from the standard parametric specification in the empirical option pricing literature. First, the price jumps are exponentially

³The assumption, that the positive jump intensity is time-invariant is, in part, motivated by the scarce availability of deep out-of-the-money call options, which renders it difficult to draw inference regarding shifts in the shape of the far right side of the option surface.

distributed, and consequently have fat tails, while most prior studies rely on Gaussian price jumps, following Merton (1976). Second, the jumps in the factors V and U are linked deterministically to the negative price jumps, with the squared jumps impacting the factor dynamics in a manner reminiscent of discrete GARCH models. Third, and most importantly for the analysis here, the jump intensity is time-varying, and with innovations that are linked directly to the volatility and price jumps, yet its dynamics is largely decoupled from volatility. This is unlike most existing models, with the notable exception of Christoffersen et al. (2012), Gruber et al. (2015) and Li and Zinna (2017). Nonetheless, the representation (1) belongs to the affine class of models of Duffie et al. (2000).⁴ For future reference, we label our two-factor affine model the $2FU$ model.

The model is motivated by recent evidence that the negative jump intensity represents a separate risk factor which is essential in capturing the dynamics in the left side of the option surface. Moreover, that study finds the left tail factor to be a critical driver of the variation in the equity and variance risk premiums. At the same time, model (1) provides a more parsimonious representation than the preferred specification in Andersen et al. (2015b), which features two separate volatility factors. This choice reflects our desire to achieve precise identification of the relevant jump factor from the shorter, less liquid option samples encountered in this study. Appendix A.5 verifies that the extracted jump intensity factor from the current model is nearly identical to the one obtained from the more elaborate model, in the sense that the two series are close to proportional over the part of the sample which is common to the two studies. That is, the jump factor remains robustly identified in the current more restricted model.

We finish this section by defining a few risk measures and the associated risk premiums that we use throughout the paper. The quadratic variation of the log-price is given by,

$$\int_t^{t+h} \sigma_s^2 ds + \int_t^{t+h} \int_{\mathbb{R}} x^2 \mu(ds, dx). \quad (3)$$

We denote its conditional risk-neutral expectation by $QV_{t,t+h}$,

$$QV_{t,t+h} = CV_{t,t+h} + JV_{t,t+h} = E_t^{\mathbb{Q}} \left[\int_t^{t+h} V_t dt \right] + E_t^{\mathbb{Q}} \left[\int_t^{t+h} \int_{\mathbb{R}} x^2 \nu_t^{\mathbb{Q}}(dx) \right], \quad (4)$$

⁴There is some option-based evidence, see, e.g., Jones (2003) and Christoffersen et al. (2010), that non-affine models work better. The fact that our key findings are driven primarily by short-maturity options should mitigate the importance of potential misspecification stemming from nonlinearities in the volatility dynamics.

where $CV_{t,t+h}$ and $JV_{t,t+h}$ are the expected diffusive and jump variation under the risk-neutral measure, respectively.⁵ The expected jump variation $JV_{t,t+h}$, can be decomposed further into terms stemming from negative and positive jumps,

$$JV_{t,t+h} = NJV_{t,t+h} + PJV_{t,t+h} = E_t^{\mathbb{Q}} \left[\int_t^{t+h} \int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx) \right] + E_t^{\mathbb{Q}} \left[\int_t^{t+h} \int_{x>0} x^2 \nu_t^{\mathbb{Q}}(dx) \right].$$

Importantly, $NJV_{t,t+h}$ is proportional to U_t , while its general level, of course, is also tied to the jump size distribution. Specifically, for the $2FU$ model, the instantaneous risk-neutral negative jump variation NJV_t (i.e., with $h = 0$) equals,

$$NJV_t = \int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx) = \frac{2}{\lambda_-^2} U_t. \quad (5)$$

In contrast, the (risk-neutral) expected positive jump variation, $PJV_{t,t+h}$, is constant and equals $c_0^+ \frac{2}{\lambda_+^2} h$. Finally, because of the presence of the jumps in V_t , $CV_{t,t+h}$ is, in general, an affine function of both V_t and U_t . However, we document later on that, empirically, almost all variation in $CV_{t,t+h}$ is due to V_t .

The above measures can also be computed under the equivalent statistical measure (\mathbb{P}), which leads to the definition of the variance risk premium as,

$$VRP_{t,t+h} = QV_{t,t+h} - QV_{t,t+h}^{\mathbb{P}},$$

and the negative jump risk premium as,

$$NJRP_{t,t+h} = NJV_{t,t+h} - NJV_{t,t+h}^{\mathbb{P}}.$$

Our computation of the expected return variation measures under the statistical \mathbb{P} measure follows standard reduced-form procedures for forecasting of volatility and jump risks. That is, we do not specify a full parametric model for X under \mathbb{P} . Instead, we use a reduced-form model to generate the necessary conditional \mathbb{P} expectations for the calculation of the above risk premiums, exploiting successful techniques adapted from the extensive prior work on this topic. Appendix A.3 provides a detailed description of this computation.

⁵For notational simplicity, we often do not include a superscript \mathbb{Q} to indicate that a specific quantity is obtained under the risk-neutral measure. For example, the expected diffusive return variation under the risk-neutral measure is denoted $CV_{t,t+h}$ rather than $CV_{t,t+h}^{\mathbb{Q}}$. This should not cause confusion as, whenever we refer to expectations under the actual or statistical probability measure, the relevant quantities will carry a superscript \mathbb{P} .

3 Data

We exploit equity-index option data for the U.S. and a number of European indices. These are supplemented by high-frequency return and futures data for the underlying equity indices.

3.1 Equity-Index Option Data

Our sample from the OptionMetrics Ivy DB Global Indices covers January 2007 – December 2014, except for the Spanish data (IBEX), which are available only from May 18, 2007 - December 31, 2014. The OptionMetrics database collects historical prices from listed index option markets worldwide. Detailed information regarding the various equity-index option contracts and the zero curve for the relevant currencies are also provided.

We obtain data for seven indices: U.S. (SPX), Europe (SX5E), Germany (DAX), Switzerland (SMI), United Kingdom (UKX), Italy (MIB), Spain (IBEX). For each index, Table A1 of the Appendix provides the exchange trading hours, which we use to align the observations with the underlying high-frequency index returns, along with a number of contractual details. Given the novelty of the database, we devote particular attention to filtering the data. For each contract at any given time, either the last trade price or the exchange settlement price is reported. While it is impossible to distinguish the two, the vendor notes that 98% of the data represent settlement prices and only 2% reflect trade prices, with some variability depending of the specific exchange.

We create the final sample through the following steps. First, for each option maturity, we compute the corresponding interest rate by interpolating the zero curve for the given country. Second, we compute the implied forward price of the underlying index using put-call parity. For this purpose we retain only option cross sections with at least 5 put-call contracts with the same strike price, and then extract the futures price exploiting the full set of option pairs with the same strike. Third, we apply a few filters to ensure that the prices are reliable. We only use options with a tenor below one year, as longer maturity contracts tend to be illiquid. However, contrary to the prior literature, we include very short-maturity options in our analysis. This is due to the recent successful introduction of short-dated options by several exchanges worldwide. These options are particularly informative regarding the current state of the return dynamics, see, e.g., Andersen et al. (2017) for details on the weekly S&P 500 options. Finally, we only retain options whose prices are at least threefold the minimum tick size.

	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
Number of Options	623	327	477	349	232	164	232
Number of Maturities	11	9	9	6	6	6	6
Q05 Moneyness	-6.24	-5.21	-5.34	-5.00	-4.05	-4.46	-2.37
Q95 Moneyness	1.74	1.80	1.80	2.13	1.53	2.02	1.34
Min. Maturity	11	9	9	20	21	15	21
Max. Maturity	329	316	316	316	316	316	316

Table 1: **Summary Statistics for the Equity-Index Option Panels.** The table reports summary statistics for the number of option contracts (first row), number of available maturities (second row), the 5th and 95th moneyness percentile (third and fourth rows), and the shortest and longest available maturities in the last two rows. Moneyness is defined as $\log(K/F)/(IV_{ATM}\sqrt{\tau})$, where K is the strike price, F is the underlying forward price, IV_{ATM} is the at-the-money implied volatility, and τ is the maturity of the option contract. The numbers represent the median of the daily values.

Table 1 reports on the characteristics of the different option markets. The U.S. has the largest market, both in terms of the number of observations and available maturities. In Europe, the DAX is the most liquid market, while the Italian market features the lowest number of option observations. Typically, every market has more than 25 options available per cross section. The moneyness range, measured in standard deviations around the underlying future price, covers, on average, the interval $[-5, 2]$. This enables robust identification of the spot volatility, which loads on at-the-money options (i.e., moneyness close to zero), and the negative jump intensity, which loads on deep out-of-the-money put options (i.e., moneyness < 0). In terms of tenor, the U.S., German, and Euro indices offer more observations for short maturity options; a consequence of the increased popularity of the so-called weekly options in these markets. The remaining markets display a more traditional pattern in the trading of option, with a median short maturity of around 15 days.

For all indices, estimation exploits the full range of data, but in the figures we typically display the series for the overlapping sample May 18, 2007 - December 31, 2014, only.

3.2 High-Frequency Equity-Index Futures Data

We obtain intra-day observations on the futures written on the underlying equity indices from Tick-Data. We extract the futures price each minute, but our (annualized) realized variation measures are based on five-minute returns, striking a balance between the number of observations and the

extent of market microstructure noise. We compute the daily realized return variation (RV), which is constructed from a measure of the total quadratic variation of the log-price over the trading day. We further split this measure into: (1) the truncated variation (TV), capturing the variation due to the diffusive returns, and (2) the jump variation (JV), reflecting the variation stemming from jumps. We also compute the negative and positive jump variations (NJV and PJV), indicating the jump variation due to negative and positive jumps. The measures are obtained following the procedure of Bollerslev and Todorov (2011); see Section A.3 for details. Table A2 of the Appendix reports, for each index, the country, the associated exchange, and various contractual features. We stress that the trading hours are not fully synchronized and are of different duration across the exchanges. In particular, the U.S. trading hours overlap with the European ones by about only two hours per day.⁶

	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
\sqrt{RV}	17.57	26.87	24.56	19.32	20.29	22.39	22.30
\sqrt{TV}	16.22	23.25	20.91	16.24	18.37	20.28	19.97
JV/RV	0.15	0.25	0.28	0.29	0.18	0.18	0.20
LJV/JV	0.50	0.54	0.55	0.56	0.54	0.56	0.54
Mean Log-Return (%)	6.83	0.79	8.18	3.05	4.42	-5.52	1.41
Std Log-Return (%)	22.30	24.76	23.67	19.31	20.87	27.74	26.51

Table 2: **Summary Statistics for the Equity-Index Futures Series.** The table reports summary statistics for the realized return variation (RV), truncated return variation (TV), jump return variation (JV), negative jump return variation (LJV), and for the daily log-returns of each index. The numbers are annualized and given in percentage form, except for the ratios in rows three and four. The various return variation measures are computed from log-returns within the trading day, using the procedure detailed in the Appendix, then averaged and scaled by 252 to represent the trading days in one calendar year. We report the square-root of these measures, so they represent annualized standard deviation units. The daily mean and standard deviation in rows five and six are computed from close-to-close index return.

Table 2 summarizes the basic features of the individual index return series. The mean returns (fifth row) signify the dramatic discrepancy in performance across the sample period, with the

⁶The trading hours are slightly ambiguous, as electronic trading takes place outside the stated interval. For example, the S&P 500 e-mini futures trade almost 24 hours on the GLOBEX platform, while the table refers to the most active period, when pit trading is also in progress. The table follows the conventions adopted by TickData.

DAX futures earning about 8.2% and the MIB futures losing 5.5% per year. The return variation measures display more parity, but the eurozone currencies are consistently the most volatile. For the U.S., Spain and Italy, the overnight returns add a substantial amount of additional variation. Nonetheless, in all cases, the Swiss, British and U.S. indices have lowest realized return variation, both based on the intraday RV measure (row one) and the daily close-to-close measures (row six). This is consistent with the eurozone indices being more sensitive to the turmoil associated with the sovereign debt crises over this period. Finally, we note that the jumps constitute a non-trivial component of the overall trading day return variation (row three), ranging from around 15% to close to 30%, with only a marginally larger impact stemming from negative jumps (row four).

4 Estimation Procedure

We follow the estimation and inference procedures in Andersen et al. (2015a) for recovering the parameters and latent factor realizations of model (1). The option prices are converted into corresponding Black-Scholes implied volatilities (BSIV), i.e., any out-of-the-money (OTM) option price observed at time t with tenor τ (measured in years) and log moneyness $k = \log(K/F_{t,t+\tau})$ is represented by the BSIV, $\kappa_{t,k,\tau}$. For a given state vector, $\mathbf{S}_t = (V_t, U_t)$, and risk-neutral parameter vector θ , the corresponding model-implied BSIV is denoted $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$. Estimation of the parameter vector and the period-by-period realization of the state vector now proceeds by minimizing the distance between observed and model-implied BSIV in a metric that also penalizes for discrepancies between the inferred spot volatilities and those estimated (in a model-free way) from high-frequency observations on the underlying asset, $\sqrt{\widehat{V}_t^n}$.⁷ The imposition of (statistical) equality between the spot volatility estimated from the actual and risk-neutral measure reflects an underlying no-arbitrage condition, which must be satisfied for the option pricing paradigm to be valid.

To formally specify the estimation criterion, we require some notation. We let $t = 1, \dots, T$, denote the dates for which we observe the option prices at the end of trading. We focus on OTM options, with $k \leq 0$ indicating OTM puts and $k > 0$ OTM calls. Due to put-call parity, there is obviously no loss of information from using only OTM options in the estimation.

We obtain point estimates for the parameter vector θ and the period-by-period state vector

⁷We obtain the spot volatility via the so-called truncated realized volatility estimator, using five-minute returns over a three-hour window prior to the close of the trading day, see Section A.3 of the Appendix for further details.

$\mathbf{S}_t = (V_t, U_t)$ from the following optimization problem, see Andersen et al. (2015a),

$$\left(\{\hat{\mathbf{S}}_t\}_{t=1}^T, \hat{\theta}\right) = \underset{\{\mathbf{S}_t\}_{t=1}^T, \theta}{\operatorname{argmin}} \sum_{t=1}^T \left\{ \sum_{\tau_j, k_j} \frac{(\kappa_{t, k_j, \tau_j} - \kappa_{k_j, \tau_j}(\mathbf{S}_t, \theta))^2}{N_t} + \frac{0.05}{N_t} \frac{(\sqrt{\hat{V}_t^n} - \sqrt{V_t})^2}{\hat{V}_t^n/2} \right\}. \quad (6)$$

To reduce the computational burden, we only estimate the system for options sampled on Wednesday or, if this date is missing, the following trading day. The critical feature ensuring good identification of the parameters is to obtain observations across heterogeneous constellations of the option surface. We achieve this by exploiting all options that pass our filter. The shape of the surface varies dramatically in the early and late years, when the market is fairly quiet, relative to the periods associated with the financial and European debt crises. Once the parameter vector and the state variable realizations for those Wednesdays have been obtained, it is straightforward to “filter” the state variables for the remaining trading days, exploiting the estimated parameters and the criterion (6). Thus, we have daily estimates for the state realizations available, even if full-fledged estimation is performed only for weekly data.

We emphasize that the estimation procedure is devoid of parametric assumptions concerning the underlying equity-index returns. We only impose the no-arbitrage condition that implies equality between the spot volatility under the risk-neutral and objective measure, while allowing for (statistical) deviations that reflect measurement errors for both the option pricing model implied volatility estimator and the nonparametric high-frequency return based spot volatility estimator.

5 Option Factors

Given our limited sample period and the lower number of observations available for some of the European indices, we do not aim for a perfect option pricing model, but rather seek a specification that captures the salient features across all indices in a robust manner. Thus, our representation involves only a single volatility factor and the dynamics of the negative jump intensity factor is pared down relative to Andersen et al. (2015b). This largely eliminates instances where separate identification of the factors is troublesome and contributes to a sharp separation of the volatility and jump features for all the indices, ensuring that the cross-country comparisons are meaningful.

Of course, one may still be concerned that an overly simplified specification of the volatility dynamics may distort the inference regarding the jump factor. Since the separation of the return

variation into a (diffusive) stochastic volatility factor and a (left) jump factor is critical for the interpretation of our empirical findings, we illustrate the effect of operating with a one- versus two-factor representation of the the stochastic volatility component in Section A.5. We exploit the fact that Andersen et al. (2015b) estimate an extended jump-diffusive two-factor version of model (1) for S&P 500 options, which provides a good fit to the option prices as well as the volatility series, clearly outperforming more standard and equally heavily parameterized representations. Comparing the extracted series for the left jump factor across the two models, we observe a remarkable degree of coherence, with a correlation close to 99%. This is achieved in spite of the model parameters being estimated over different, only partially overlapping, sample periods, and exploiting a different set of underlying options.

In summary, we find that the main tenets of our analysis for the S&P 500 index are robust to a more refined modeling of the volatility dynamics, and the qualitative conclusions are not affected by switching to the more involved model. A detailed account of our estimation results for the S&P 500 index as well as the European indices is provided in Section A.6 of the Appendix.

5.1 Country-by-Country Factor Realizations

We now explore the implied factors for the various countries extracted from the equity-index option markets based on our $2FU$ model (1) and the estimation procedure outlined in Section 4. We first compare the implied spot variance of the European indices enumerated in Section 3, in each case exploiting the spot variance factor for the S&P 500 options as a benchmark.

Figure 2 plots the extracted volatilities from the Euro ESTOXX (Eurozone), DAX (Germany), SMI (Switzerland), FTSE 100 (U.K.), FTSE MIB (Italy), and IBEX 35 (Spain) markets. The most striking feature is the extraordinary close association between many of these factors and the S&P 500 spot volatility, not just in terms of correlation, but also level. For example, the U.K. volatility is barely distinguishable from the S&P 500 factor throughout the sample, while the Swiss factor deviates visibly from the S&P 500 factor only during a few episodes following the Swiss franc-euro exchange rate cap implemented in September 2011. In the former case, the volatility correlation is about 98% and in the latter around 95%.

In contrast, notable discrepancies between the S&P 500 and German DAX indices emerge during the second phase of the European sovereign debt crisis, when the DAX volatility spikes relatively

more, and a positive gap remains from that point onward, albeit to a varying degree. For the broader euro-zone ESTOXX index, the same effect is clearly visible and originates with the initial phase of the European debt crisis. Thus, while the volatility patterns were very similar for all the primary equity indices in North America and Europe through the financial crisis, the response to the sovereign debt crisis is heterogeneous, with the relative impact corresponding nicely to our perception regarding the sensitivity of the respective economies to the European crisis. This is especially striking for the Italian and Spanish indices, as both react very strongly to the crisis events, but with different amplitudes across the main episodes. The volatility levels for these two Southern European indices attain a plateau well above the others ever since the sovereign debt problems surfaced in early 2010. This systematic divergence over the second half of our sample lowers the volatility correlations for MIB and IBEX with S&P 500 to 0.75 and 0.77, respectively.

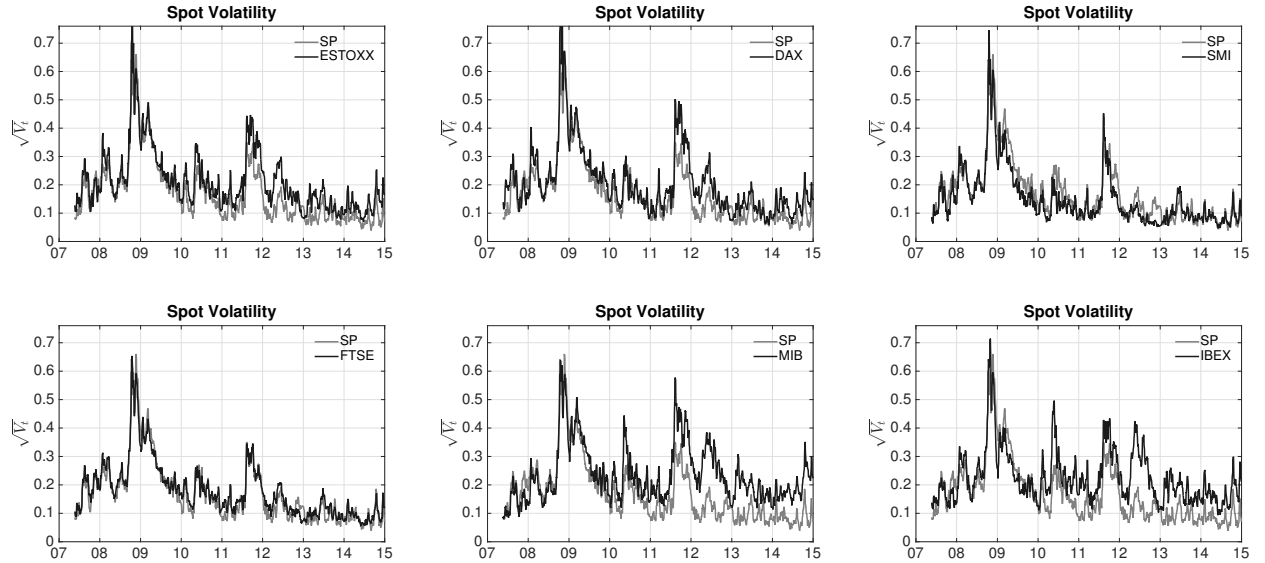


Figure 2: **Volatility Factor Comparison.** For each option-implied spot variance factor, obtained at the close of the trading day, we report the trailing five-day moving average of $\sqrt{V_t}$. The series are given in decimals and refer to annualized values.

Next, Figure 3 depicts the option-implied (risk-neutral) negative jump variation, NJV_t , for each index along with the corresponding quantity for the S&P 500. As noted in Section 2, the risk-neutral negative jump variation in the $2FU$ model, $\frac{2}{\lambda_-^2}U_t$, is proportional to the negative jump intensity factor, U . Consequently, the relative variation in the series reflects the corresponding variation in the extracted jump intensity factor for the individual indices.

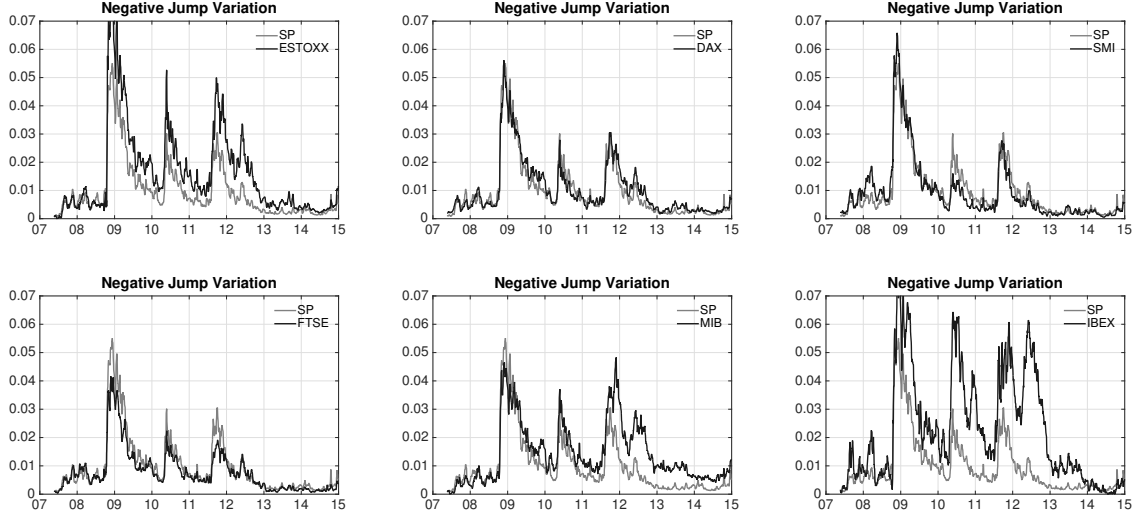


Figure 3: **U Factor Comparison.** For each option-implied negative jump intensity factor, we report the trailing five-day moving average of the implied risk-neutral jump variation $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx) = NJV_t \equiv \frac{2}{\lambda_-^2} U_t$. The series are given in decimals and refer to annualized values.

At first glance, the general pattern is similar to the one observed for the volatility factors. This is due to the fact that the volatility and jump factors are quite highly correlated for all countries. Nevertheless, the relative size of the spikes in the jump intensity versus volatility varies substantially, with the sovereign debt crisis inducing a stronger surge in the jump intensity relative to (diffusive) volatility. As for the volatility factors, the U.K. and U.S. series evolve in near unison and display a correlation above 98%, even if the British jump variation is slightly lower throughout. Similarly, the jump intensity factor for the Swiss index correlates strongly with the S&P 500 factor, although the Swiss factor tends to be above the one in the U.S., both before and during the financial crisis, and then below after the summer of 2009.

We also observe a strong coherence between the S&P 500, ESTOXX and DAX series up through the financial crisis and then a relative elevation in the latter two from the summer of 2009 and onwards, with the effect being notably more pronounced for ESTOXX than DAX, again suggesting a smaller exposure of Germany to the debt crisis than for the broader euro-zone. As before, however, the most striking contrast occurs for the two Southern European indices. The Italian jump variation spikes to a level corresponding to the financial crisis during the latter part of 2011, and the Spanish one is exceptionally highly elevated during several phases of the sovereign debt crisis. For these two countries, the jump intensities convey a very different impression of the severity of the debt crisis,

both relative to the other countries and to the corresponding volatility factors. This is perhaps not surprising given the widespread speculation at the time that either country might be forced to abandon the euro currency. In summary, our decomposition of the primary risk factors documents a substantially larger increase in return volatility for these Southern European indices along with a further amplification of the negative jump risk, especially for Spain.

The coherence across the volatility and jump variation series as well as the striking, but economically interpretable, discrepancies observed during crisis episodes add credence to the robustness of our methodology in extracting the salient features from the option surfaces.

5.2 Option Factors as Predictors of Future Risk and Risk Premiums

In affine models, the state variables governing the risk-neutral return dynamics are generally tied to the underlying return dynamics and the associated compensation for risk. Hence, we now explore the ability of the two option-implied factors, the spot variance and the left jump intensity, to forecast the (realized) return variation – defined as the sum of the squared high-frequency equity-index futures returns – and the equity excess return. The former signifies whether the factors are associated with market-wide risk, as captured by the ex-post realized equity volatility and jump activity, while the latter speaks to their forecast power for future equity returns and, more generally, the equity risk premium.

A common concern with predictive regressions is the potential look-ahead bias embedded in the regressors. This is not a major issue in our setting. The option factors are extracted daily, using only the option prices and the high-frequency volatility measure for that given day, along with the parameter estimates, obtained strictly from the varying shape of the option surface. As such, the daily or weekly index returns play no role in our estimation and factor extraction procedures.⁸

We further note that the option factors, depicted in Figures 2 and 3, are quite persistent, yet not to an extent that render them near-unit-root processes. This mitigates concerns regarding inference that may arise from excessive persistence in the (option factor) regressors.⁹

⁸In results, available upon request, we document that the option-implied factors are very similar, and provide qualitatively identical forecast performance, whether the parameters are estimated across the full sample or from an initial year of data only. In the latter case, the parameter vector is fixed at the point estimate based on the early parts of the sample and then used for extracting factors and pricing options over the remaining (out-of-sample) period.

⁹For a discussion of the problems induced by treating the regressors as near non-stationary when, in fact, they are strictly stationary, see Phillips (2014).

5.2.1 Predicting Equity Risk

What do the option-implied factors tell us about the risk characteristics of the underlying equity-index? To explore this issue, we regress a measure of the future realized return variation, $RV_{t,t+h}$, over the time interval $[t, t+h]$, on the option-implied state variables. The $RV_{t,t+h}$ measure is constructed from high-frequency intraday observations on the equity-index futures augmented with the squared overnight returns. The high-frequency data afford accurate measurement of the ex-post return variability, stemming from diffusive volatility and jumps, so they provide good proxies for the future risks associated with exposure to the market index, see, e.g., Andersen et al. (2003).

Since the two state variables are highly correlated for all the indices, for ease of interpretation, we supplement the spot variance factor, V , with the component of the negative jump variation factor that is orthogonal to the spot variance, denoted NJV^\perp , as the second explanatory state variable. This approach ascribes all predictive power stemming from joint variation in the state variables to the traditional spot variance factor, while the residual variation in the left tail factor captures only the explanatory power of the regressors that is unrelated to the variance factor, i.e., it reflects purely the incremental information in the tail factor.

Given our limited sample period, we run the predictive regression on a weekly basis, forecasting from 1 to 28 weeks, or roughly 6 months, into the future. Due to the short sample period and the varying liquidity for some of our index options, the results can be sensitive to outliers. The more extreme observations may be genuine, but substantial measurement errors can also arise from a variety of sources, including periodic illiquidity in the index option markets, large option bid-ask spreads during turbulent events, data errors, non-synchronous observations, large standard errors for the nonparametric high-frequency volatility estimators on days with elevated volatility, and potential model misspecification on trading days with unusual market stress. Such errors may induce poor identification of the factor realizations. Hence, for robustness, we winsorize all explanatory variables in the subsequent predictive regressions at the 98 percent level, thus limiting the influence of the 1% extreme negative and positive observations. Importantly, we do not alter the corresponding daily return variation or excess return from the cumulative measures appearing as left-hand-side variables, i.e., all extreme volatility and return realizations are included in the multi-horizon excess returns and return variation measures. The inference is based on Newey-West robust standard errors with a lag length equaling twice the number of weeks within the forecast

horizon, and they also account for the estimation errors in any potential first-stage regression used to normalize the regressors.

The regressions take the form,

$$RV_{t,t+h} = k_{0,h} + k_{v,h} \cdot V_t + k_{u,h} \cdot NJV_t^\perp + \epsilon_{t,h}, \quad (7)$$

where we reiterate that the expected monthly negative jump variation under the risk-neutral measure, NJV_t^\perp , is directly proportional to the component of the state variable U_t that is orthogonal to the spot variance, as implied by equation (5). The time series of NJV_t^\perp for the different indices is reported in Figures A3 and A4.

The left panels of Figure 4 reveal, for all our equity indices, that the part of the left jump intensity factor orthogonal to spot volatility has no explanatory power for the ex-post realized return variation. Instead, all predictor power is concentrated in the implied spot variance which, of course, is well known to be a powerful predictor of short-term return volatility. The right panels show that the explained variation is very high and qualitatively similar across all indices. The absence of any auxiliary predictive power in the jump factor is striking. It signifies a stark disconnect between the residual movements in the left side of the implied volatility surface (the part not spanned by volatility) and the riskiness, or variation, of the future returns.

5.2.2 Predicting Equity Excess Returns

An important implication of the findings in Section 5.2.1 is that the orthogonal component of the negative jump variation, NJV^\perp , will not be recognized as a factor driving any facet of the risk dynamics in standard time series models estimated from the returns of the underlying asset. As such, many traditional approaches, that extract the relevant risk factors from the underlying return dynamics, will fail to recognize NJV^\perp as a relevant factor for option pricing. Indeed, the question is whether this option tail factor is a purely idiosyncratic feature of the pricing of OTM put options, unrelated to both risk and risk pricing across the broader asset markets. That is, it may simply represent market segmentation that arises from frictions and clientele effects in the trading of derivatives securities.

The above hypothesis would imply that the tail factor should possess no auxiliary explanatory power for the pricing of equity risk. We explore this conjecture by running predictive return




Figure 4: **Predictive Regressions for Return Variation.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

regressions analogous to those for the return risk in Section 5.2.1. Once again, we rely on mild winsorizing of the regressors, and we continue to ascribe all explanatory power, that stem from joint variation in the state variables, to the traditional spot variance factor.

Hence, as for equation (7), for each index and $t = 1, \dots, T - h$, our regressions take the form,

$$r_{t,t+h} = \log(X_{t+h}) - \log(X_t) = c_{0,h} + c_{v,h} \cdot V_t + c_{u,h} \cdot NJV_t^\perp + \epsilon_{t,h}. \quad (8)$$

The predictive regressions (8), summarized in Figure 5, deliver two major results. First, they indicate a substantial degree of return predictability at intermediate horizons of two to six months. Second, they consistently point to the orthogonal left tail factor, not the volatility, as the relevant indicator of equity risk pricing. As such, we have a complete reversal of the roles of our two factors: the significant explanatory variable is now the tail factor, while the volatility factor is largely irrelevant for the future returns.

For example, the left panel in the first and second row of Figure 5 plot the t-statistics for the SP and ESTOXX regression slopes in equation (8), while the right panels display the corresponding R^2 statistics. For both indices the predictive power is low at high frequencies, but then rises steadily with the horizon until about five months.¹⁰ Within the four-month mark, the R^2 surpasses 10%, and it exceeds 15% after six months. The increasing forecast power for longer return horizons is consistent with the hypothesis of a time-varying and persistent equity risk premium, combined with a second mildly persistent and return component, that is correlated with the regressors, see, e.g., the discussion in Stambaugh (1999) and Sizova (2016). The interpretation is that, at the weekly horizon, the largely unpredictable, noisy short-term component dominates whereas, for longer holding periods, the predictable return component emerges.

The truly striking feature of Figure 5 is, however, as noted above, that the explanatory power stems almost exclusively from the jump intensity factor, as the variance factor is insignificant across all horizons for the S&P 500 and ESTOXX indices. Thus, the commonly employed volatility factor has no discernible relationship with the equity risk premium, while unrelated variation in the left side of the option surface is indicative of systematic shifts in the pricing of equity risk.

¹⁰This qualitative pattern is familiar, as it is observed for other return predictor variables as well, including the dividend-price ratio and moving averages of interest rates. Yet, there is an important distinction, as these predictors are much more persistent, and they attain significance only at much longer multi-year horizons.

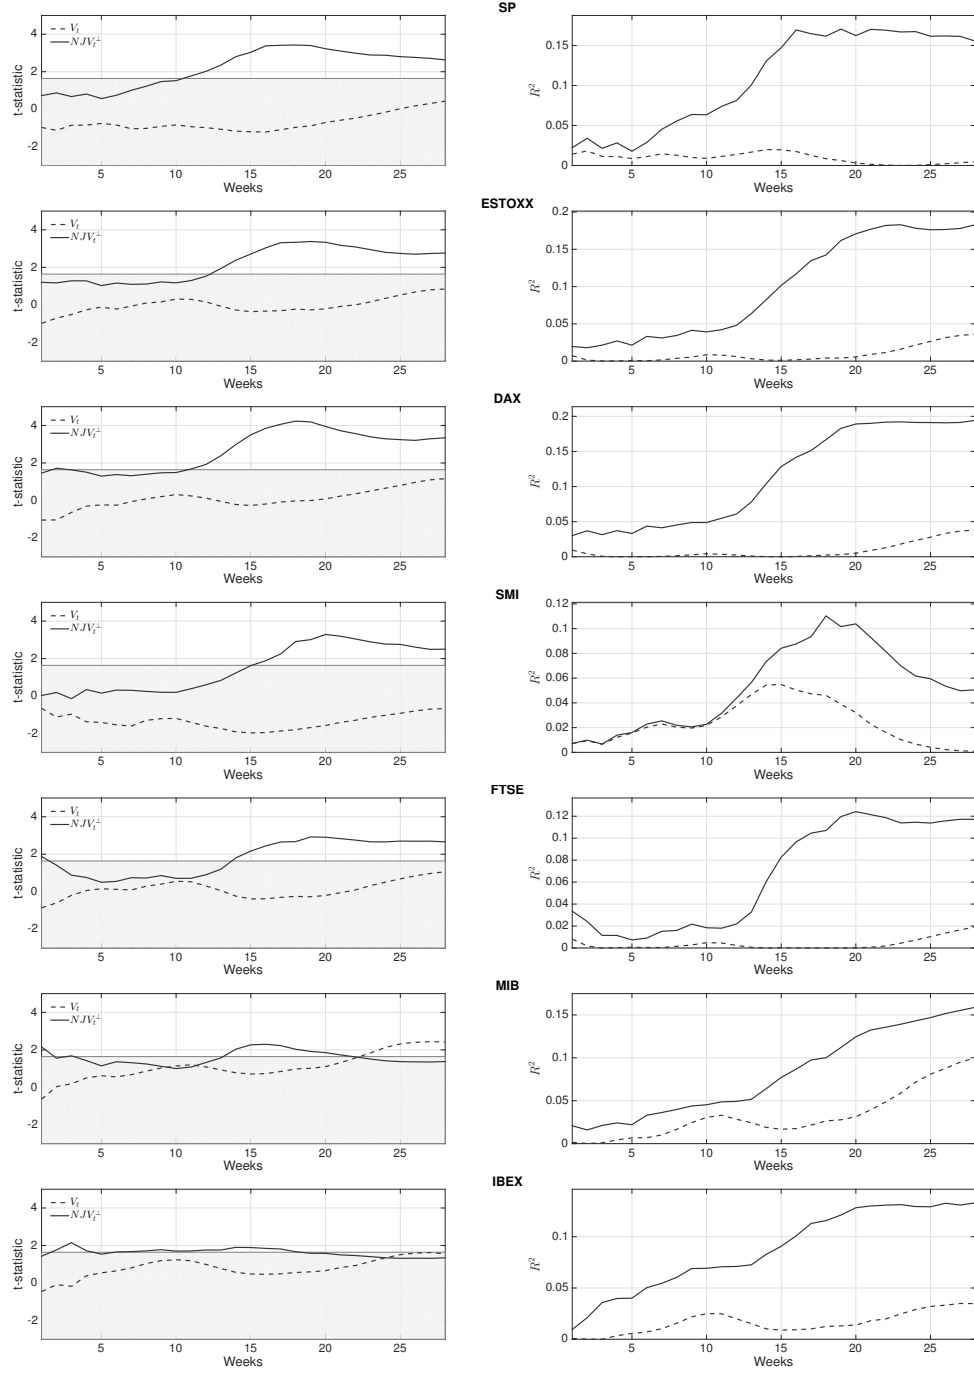


Figure 5: **Predictive Regressions for Excess Returns.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

Turning to the remaining indices in Figure 5, we observe qualitatively similar features across the board. The only noteworthy differences appear for the two Southern European indices MIB and IBEX, which display a slightly lower degree of significance for the residual left tail factor at the (one-sided) 5% level and some explanatory power for the volatility factor at the longest horizons. Since the evidence for predictability beyond five-six months must be viewed somewhat skeptically given the short sample, this is not a surprising finding for two countries subject to extreme dislocations in the market during both the financial crisis and the subsequent sovereign debt crises.

Moreover, we note that the degree of explained variation consistently falls between 12-20% at the longer horizons, excluding the Swiss index. For Switzerland, the sudden implementation of an exchange rate cap of the Swiss franc versus the euro on September 6th, 2011, represented a major shock to the Swiss equity market, rendering the explanatory power lower than for the other indices, even if the results are not qualitatively different. Specifically, the sharp depreciation of the franc at the introduction of the cap was accompanied by a large positive jump in the (franc denominated) index. Since this intervention was unprecedented, and certainly unexpected, the sharp appreciation of the local index was not reflected *a priori* in the option surface.¹¹

We conclude that the variance factor, effectively, is bereft of explanatory power in these regressions. In contrast, the orthogonalized tail factor provides robust predictive power for the excess returns across all indices. This finding is remarkable given the huge discrepancy in the realized index returns over the sample, and the diverse exposures they exhibit vis-a-vis the European debt crises. We also note that the insignificant volatility factor is in line with an extensive time series literature, which has failed to generate consistent evidence that the equity-index return volatility predicts future equity returns, see, e.g., French et al. (1987) and Glosten et al. (1993) for early references. Importantly, our evidence suggests that the option surface does embed critical information for future market returns, but it is contained within factors that are unspanned by volatility.

¹¹Furthermore, since the appreciation of the index was smaller than the simultaneous devaluation of the Swiss franc, the equity market performance, *per se*, is a poor guide to the return performance viewed from a global perspective. These concerns are even more relevant for the discussion of return predictability associated with the variance risk premium in the following section. The dismantling of the exchange rate cap in 2015 also took the markets by surprise, but it occurred after the end of our sample period.

6 Option Factors and The Pricing of Risk

The risk-neutral dynamics and factor realizations, extracted from the individual equity-index option panels, convey information about the expected future return variation, the underlying sources of risk, and the pricing of those risks. In Section 5.2, the focus was exclusively on the option factors and their information content. In this section, we explore more directly the pricing of different sources of risk, their interaction, and their connection to the extracted option factors.

6.1 The Variance Risk Premium

We start with the variance risk premium, which has been the subject of much interest in the recent literature. It is well-established empirically that option-implied volatility measures possess strong predictive power for the future return variation. Figure 4 confirms that this is, indeed, the case for all our equity indices as well. Moreover, these equity risk measures display pronounced time variation and a significant degree of persistence, as evident for the volatility series in Figure 2. These features provide, qualitatively, a rationale for a (negative) premium on payoffs linked to the future level of volatility, reflecting the compensation investors demand for bearing (systematic) variance risk. Formally, this manifests itself in a gap between the conditional risk-neutral and statistical expectation of the return variation, corresponding to our definition of the variance risk premium, VRP , in Section 2.

In terms of practical measurement, the risk-neutral expectation of the quadratic return variation is readily obtained nonparametrically using a portfolio of close-to-maturity options, and it corresponds closely to the value of the VIX index at the corresponding maturity.¹² We follow the procedure of Carr and Wu (2009) for computing the risk-neutral expected return variation for the 30-day horizon, outlined in Appendix A.4. For the corresponding statistical measure, we rely on standard forecasting techniques for the expected return variation over a 30-day horizon exploiting high-frequency data, as outlined in Section A.3 of the Appendix.¹³

The extant literature has consistently found large negative variance risk premiums for equity indices. For the indices in our sample, the average estimate for the VRP ranges from -1.6% to

¹²See, for example, Bakshi and Madan (2000) and Carr and Wu (2009) for the construction of model-free measures from options portfolios. Andersen et al. (2015) document that the VIX index provides a close approximation to the risk-neutral expectation of the future return variation, even in the presence of jumps in the return generating process.

¹³An alternative is to rely on our estimated risk-neutral model for the return dynamics and the extracted option factors to construct the VIX measure. The two approaches produce very similar results.

−3.9% in annualized volatility terms.¹⁴ Given the relatively short time span, these estimates are somewhat noisy, but they are, nonetheless, entirely in line with prior evidence.

Recent studies further find that, both for U.S. and international equity indices, the country-specific variance risk premium has predictive power for that country’s future excess returns; see, e.g., Bollerslev et al. (2009) and Bollerslev et al. (2014).¹⁵ In our setting, the state vector should determine the risk-neutral expected quadratic return variation and simultaneously, as discussed above, it provides a good forecast for the future expected return variation (under the statistical, or \mathbb{P} , measure). Consequently, an affine mapping links the state vector to (a good proxy for) the variance risk premium. The fact that we obtain significant predictive power for the future excess returns from the bivariate regression involving the two (orthogonalized) factors, V and NJV^\perp , for horizons between two to six months in Figure 5 is consistent with this general line of reasoning.

Nonetheless, our findings raise critical questions regarding the association between the variance and equity risk premiums. In particular, as detailed in Section 5.2.2, V has essentially no explanatory power for the equity risk premium. Instead, the predictive power resides squarely with the pure (orthogonal) negative jump factor, NJV^\perp . This suggests that only certain specific components of the variance risk premium have explanatory power for the future equity risk premium. This fact has important implications for our understanding of the pricing of equity risk. Hence, in the following subsections we explore the impact of the left tail factor, NJV^\perp , in determining the variance and left jump risk premiums, and ultimately in providing information regarding the dynamics of the equity risk premium.

6.2 The Role of the Left Tail Option Factor

For clarity, we organize the exposition around a stark composite null hypothesis. Our prior findings suggest – but do not conclusively document – that the pure tail factor, NJV^\perp , solely reflects the investors attitude towards abrupt downside risk exposure, but carries no information regarding the actual future expected negative jump risk. That is, we conjecture that this factor constitutes a pure risk premium with no relation to actual future equity risks. Moreover, we stipulate that this

¹⁴The realized volatility used to compute the VRP corresponds to the one reported in Table 2 augmented with the the overnight (close-to-open) squared returns to ensure compatibility with the option-implied volatilities.

¹⁵This issue also relates to the broader literature on predictability for international equity indices, e.g., Harvey (1991), Bekaert and Hodrick (1992), Campbell and Hamao (1992), Ferson and Harvey (1993) and Hjalmarrsson (2010). In addition, see Bakshi et al. (2011) for the predictive ability of other option-based volatility measures.

tail factor provides a good proxy for the overall negative jump risk premium, and that it is the only element of the variance risk premium with predictive power for the equity risk premium.

To shed further light on these assertions, it is convenient to exploit the decomposition, $NJV = NJV^\parallel + NJV^\perp$, where NJV^\parallel denotes the part of NJV spanned (linearly) by V . We then have the following decomposition of the variance risk premium,

$$\begin{aligned} VRP_{t,t+h} &= NJRP_{t,t+h} + CVRP_{t,t+h} + PJRP_{t,t+h} \\ &= NJV_{t,t+h}^\perp + \left(NJV_{t,t+h}^\parallel - E_t^\mathbb{P}[NJV_{t,t+h}] \right) + \left(CV_{t,t+h} + PJV_{t,t+h} - E_t^\mathbb{P}[CV_{t,t+h} + PJV_{t,t+h}] \right), \end{aligned} \quad (9)$$

where $CVRP_{t,t+h} = E_t^\mathbb{Q}[\int_t^{t+h} V_u du] - E_t^\mathbb{P}[\int_t^{t+h} V_u du]$, while $NJRP_{t,t+h}$ is defined in Section 2, and $PJRP_{t,t+h}$ denotes the analogous risk premium, but for the positive jump variation, and, finally, the various risk-neutral return variation measures are introduced in Section 2.

Our null hypothesis imposes restrictions on the interaction among the terms in equation (9). In particular, the assertion that NJV^\perp is a pure jump risk premium component implies that it is uncorrelated with both of the objective expectation terms in the second line, i.e., $E_t^\mathbb{P}[NJV_{t,t+h}]$ and $E_t^\mathbb{P}[CV_{t,t+h} + PJV_{t,t+h}]$. The additional conjecture, that variation in NJV^\perp provides a good approximation to the overall variation in the negative jump risk premium, requires the second term to display a substantially lower degree of variation than NJV^\perp . Of course, since the second term takes the form of a premium, only the gap between the risk-neutral and future expected variation measure is included, which dampens the fluctuations of this term relative to NJV^\perp . Finally, NJV^\perp should be uncorrelated with the full expression in the third term, encompassing the risk compensation for continuous and positive jump return variation.

In addition, we obtain a rank-ordering regarding the predictability of future equity returns. The null hypothesis implies that the pure jump factor, NJV^\perp , is the true source of predictability. Hence, if we consider the standard predictive regressions setting,

$$r_{t,t+h} = d_{0,h} + d_{v,h} \cdot Y_t + \epsilon_{t,h}^Y, \quad (10)$$

where Y_t signifies the predictor variable, we obtain the following empirical implications. First, the best forecast variable for the equity premium across all option factors and risk premiums in equation (10), in terms of the degree of explained variation, R^2 , and significance of the slope coefficient, $d_{v,h}$, is the pure jump factor, $Y_t = NJV_t^\perp$. Second best is the negative jump risk premium, $Y_t = NJRP_{t,t+h}$. It comprises the sum of the first two terms in the bottom line of

equation (9). Since the second component, per assumption, has no significant predictive power for future equity returns, it merely adds noise to the predictor, generating an errors-in-variable problem, thus lowering the signal-to-noise ratio and weakening the forecast performance. Third, adding the third term from the second line of equation (9) to obtain the full variance risk premium, i.e., $Y_t = VRP_{t,t+h}$, only further aggravates the errors-in-variable problem, because the latter displays a great deal of variation over the sample but, under the null hypothesis, possesses no predictive power. Fourth, the *VIX* should perform even worse, as it includes the full risk-neutral return variation measure, adding NJV^\parallel , CV , and PJV to NJV^\perp , yet it does not include the offsets provided by the corresponding expected future objective return variation measures. Therefore, the degree of variation in NJV^\perp is only a small fraction of that for *VIX*, and the signal-to-noise ratio is now very small. Yet, the *VIX* index does embed the tail factor, and thus it may retain marginal predictive power. Finally, the spot volatility factor does not encompass the NJV^\perp measure at all, and so should fail to predict the future returns altogether. We have already documented this fact empirically in Figure 8 and it was, in fact, one of the key findings, motivating the specification of the null hypothesis that we entertain here.

6.3 The Negative Jump Risk Premium

In this section, we focus on the pure left jump factor and properties of the overall negative jump risk premium. In particular, we explore some of the assertions put forth in Section 6.2.

First, in order to establish that NJV_t^\perp is associated solely with the jump risk premium, we must demonstrate that it is not (significantly) correlated with any individual component of the expected future return variation, especially $E_t^{\mathbb{P}}[NJV_{t,t+h}]$. Even though we previously found NJV^\perp to have no predictive power for the overall return variation in equation (7), this may be due to an offset between the individual terms, or it may stem from low power arising from the dominant role of the expected continuous variation (CV) in the total return variation. If the CV component is not related to NJV^\perp , correlation with the future negative jump activity, accounting for a much smaller fraction of the actual return variation, may go undetected.

In fact, upon running an additional set of predictive regressions analogous to equation (7), but with the realized negative jump intensity as regressand, we confirm that NJV^\perp has no explanatory power for the future negative jump variation. These results are provided in Figure A2 of the

Appendix. Moreover, as expected, the identical finding holds true for the corresponding regressions with the continuous (and positive jump) return variation as regressand. We conclude that NJV^\perp is not a risk factor, but solely a risk premium factor.

Second, one may wonder whether NJV^\perp is related only to the future jump activity, or if it is associated also with the future risk-neutral continuous return variation. In theory, the latter is a distinct possibility, as the jump intensity determines the expected number of negative return jumps, and these jumps coincide with those in volatility. However, empirically, this channel is very weak. Across all indices, we find that at least 98% of the variation of $CV_{t,t+h}$ is explained by V_t , which is orthogonal to NJV^\perp . We conclude that NJV^\perp has no meaningful impact on the risk-neutral return variation, except through the negative jump variation. Furthermore, it is, effectively, unrelated to all expected future objective return variation components. This implies that we may treat NJV^\perp exclusively as a contributor to the negative jump risk premium.

Third, we assess the importance of fluctuations in NJV^\perp for the overall variation in NJV , where $NJV = NJV^\parallel + NJV^\perp$. We find that the share of variation stemming from NJV^\perp ranges from 32% for the S&P 500 index to 74% for the Spanish IBEX index, with a mean value close to 50%.¹⁶ Thus, roughly half of the movement in the option-implied tail factor, NJV , represents a pure shift in the jump risk premium, not related to the expectation regarding the future jump variation. We stress that this fraction constitutes a weak lower bound on the importance of NJV^\perp . In the second line of equation (9), the NJV^\parallel term will, to a large extent, be offset by the expected negative jump variation, $E_t^\mathbb{P}[NJV_{t,t+h}]$. That is, time-variation in the overall premium, $NJRP$, is likely dominated by NJV^\perp . We defer further evidence on this point to the following section.

We conclude that our orthogonal left jump factor constitutes a genuine component of the negative jump risk premium and is unrelated to the future expected return variation. In other words, variation in NJV^\perp translates, one-for-one, into $NJRP$. Since this pure jump factor is the only state variable with forecast power for the equity risk premium, it is now evident that the VRP predicts future returns only because of the embedded negative jump risk premium. This suggests that we should pay particular attention to the properties of $NJRP$.

Figure 6 depicts the time series of estimated negative jump risk premiums for our indices. At

¹⁶This follows readily from the R^2 statistic associated with the original univariate regression of U_t on V_t , where NJV^\perp is proportional to the regression residual. The variation shares for the remaining indices are: 59% for ESTOXX; 46% for DAX; 46% for DAX; 39% for SMI; 38% for FTSE; and 53% for MIB.

any point in time, the premium reflects the extracted jump and volatility factors, the estimates for the risk-neutral return dynamics, and the estimation procedure for the expected negative jump variation. As noted previously, the latter, i.e., the estimates for the future negative jump variation, are fairly noisy or imprecise, as jump intensities and sizes are very hard to forecast. Thus, there is some uncertainty associated with these series, but the fact that the NJV often is an order of magnitude larger than the actual objective expectation of the negative jump variation renders the pronounced systematic variation in the series reliable.

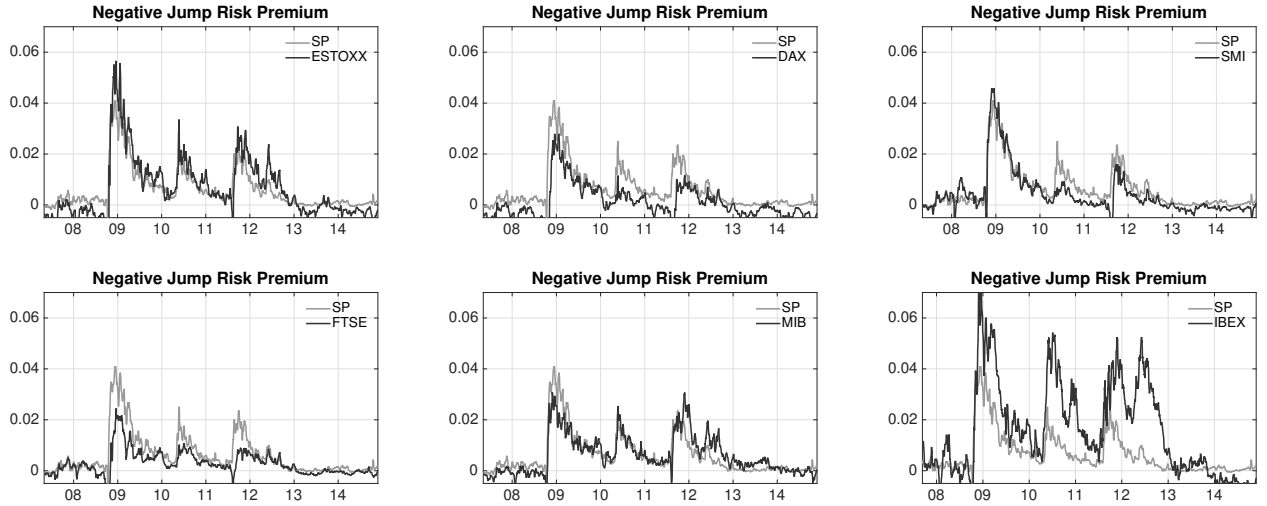


Figure 6: **Negative Jump Risk Premium.** For each index we report the negative jump risk premium as the difference between the risk-neutral and physical expectation of the negative jump variation over thirty calendar days. The series are given in decimals and refer to annualized values.

As before, the U.S. index serves as a benchmark for the series in Figure 6. We observe that, qualitatively, $NJRP$ evolves similarly to the U factor in Figure 3. Nevertheless, the distinct forecasts for the actual negative jump variation in each index generate interesting discrepancies across the two set of figures. The premiums are closely aligned throughout for the S&P 500 and ESTOXX indices while, relative to the S&P 500, they are generally lower in Germany and Britain from the financial crisis onward, lower in Switzerland after the financial crisis, and higher in Italy and Spain ever since the financial crisis, with the jump premiums in Spain becoming extraordinarily large late in the sample.

Our assertion that $NJRP$, or components thereof, possesses superior predictive power for the equity risk premium relative to the commonly adopted VRP will be explored in Section 6.5.

6.4 Interaction among Factors and Risk Premiums

Sections 6.2 and 6.3 provide evidence regarding the variability of the pure left tail factor and its impact on the variance and jump risk premiums. We now explore the interaction among the risk premiums and factors more broadly, while highlighting the unique role played by the tail factor.

Table 3 reports the sample correlation between various combinations of factors and risk premiums. The reported correlations may be biased towards zero due to imperfect procedures for forecasting the expected return variation measures entering the risk premium computations. Moreover, they may be subject to noise stemming from a few influential outliers. Nonetheless, the general coherence of the series across the full set of indices should be informative regarding the underlying relations among these quantities. This is corroborated by relative robustness of the correlations over subsamples, as demonstrated in Section A.8 of the Appendix for the periods 2007-2009 and 2010-2014, respectively. The correlations for the subsamples are generally consistent with those for the full sample period, although the smaller samples naturally do produce more erratic statistics.

	V, NJV	VRP, V	VRP, NJV^\perp	$NJRP, V$	$NJRP, NJV^\perp$	$CVRP, V$	$CVRP, NJV^\perp$
SP	0.85	0.86	0.44	0.78	0.74	0.86	0.15
ESTOXX	0.72	0.83	0.41	0.47	0.95	0.86	-0.07
DAX	0.76	0.72	0.42	0.16	0.97	0.88	0.06
SMI	0.83	0.92	0.33	0.64	0.89	0.92	0.12
FTSE	0.82	0.65	0.02	0.33	0.96	0.72	-0.28
MIB	0.80	0.89	0.15	0.68	0.84	0.86	-0.08
IBEX	0.59	0.91	0.31	0.49	0.92	0.83	-0.29

Table 3: **Factor and Risk Premium Correlations.** The table reports the sample correlation of daily observations for a combination of option factors and risk premiums over the period 2007-2014.

The first column of Table 3 reveals, as also evident from Figures 2 and 3, that the volatility and negative jump factors are highly correlated. The remaining columns in Table 3 are organized in pairs, reflecting the correlation of the two factors with a specific risk premium. The correlation of the spot variance with the variance risk premium is consistently high across indices. Clearly, whenever the environment is turbulent, the variance risk premium tends to rise. The variance risk premium

is also, but to a lesser extent, positively correlated with the pure (orthogonalized) jump factor. That is, even if the jump factor reflects the (risk-neutral) jump variation not contemporaneously accounted for by spot volatility, it is still intimately related to the variance risk premium. Of course, this partially reflects the fact that the pure jump factor accounts for a substantial part of the jump risk premium and, as such, constitutes a component of the variance risk premium.

The next pair of columns display the correlations of the option factors with the negative jump risk premium. The correlations of the negative jump risk premium with the jump factor NJV^\perp are of particular interest. The numbers are uniformly high, lending support to the hypothesis that the pure jump factor is the critical ingredient of this premium. The relatively low correlation reported for the S&P 500 index stems from a single huge (negative) deviation between the V and NJV^\perp at the onset of the financial crisis in 2008, where our option pricing model identifies the initial spike in the option prices primarily as a volatility rather than a jump intensity shock. Inspection of Table A10 in the Appendix shows that the correlation reaches 0.97 for the 2010-2014 period, fully in line with the remaining indices, which all attain a value of 0.90 or above over this subsample. In short, the pure left jump factor, NJV^\perp , is uniformly highly correlated with the jump risk premium – a particularly striking feature given the difficulty in, and noise generated by, estimating the expectations regarding the future negative jump variation. Hence, as hypothesized previously, NJV^\perp may serve as a good proxy for the negative jump risk premium.¹⁷

Finally, the last two columns of Table 3 document a high correlation between the volatility factor and the risk premium associated with the continuous return variation, while the latter is approximately uncorrelated with the pure jump factor. This observation is telling. It implies that the variance risk premium only correlates robustly with the NJV^\perp factor, as reported in columns two and three, because the former includes the jump risk premium. Once the negative jump risk premium is stripped from the variance risk premium, the compensation for diffusive risk has no apparent relation to our primary indicator for negative jump risk compensation.

¹⁷Figures A3 and A4 in the Appendix display normalized time series for the two factors along with the variance and negative jump risk premium, respectively. The close coherence between the volatility factor and variance risk premium is apparent for all indices. In contrast, for the pure left jump factor, the coherence with the negative jump risk premium is evident whenever there is a marked discrepancy in the volatility and jump risk premiums. As noted in the text, the exception is the (outlier) behavior of the S&P 500 $NJRP$ measure in late October and November of 2008, where the $NJRP$ is well aligned with the extreme spike in spot volatility.

6.5 Risk Premiums as Predictors of Future Excess Returns

The findings of Section 6.4 suggest that the variance risk premium possesses explanatory power for future equity returns solely because it incorporates the negative jump risk premium. In this section, we explore this hypothesis directly. We also relate our findings to the rank-ordering of various predictors for the equity risk premium in Section 6.2.

In Table 4, we report on predictive regressions that directly contrast the forecast power of the variance risk premium with that of the negative jump risk premium. The table refers to (constrained or unconstrained) variants of the following regression,

$$r_{t,t+h} = b_{0,h} + b_{v,h} \cdot VRP_{t,t+30} + b_{j,h} \cdot NJRP_{t,t+30} + \epsilon_{t,h}. \quad (11)$$

Our first regression seeks to verify whether the *VRP* has explanatory power, as asserted in the literature, so it imposes the constraint $b_{j,h} = 0$ in equation (11). Next, we check if the explanatory power diminishes, as we strip the *NJRP* from the *VRP*, i.e., second regression imposes the constraint $b_{v,h} = -b_{j,h}$. The third variant is simply the unconstrained regression which speaks to whether the inclusion of the *NJRP* adds auxiliary explanatory power beyond that of the *VRP*. Finally, we impose $b_{v,h} = 0$ to gauge the predictive power of the *NJRP* in isolation.

The results for forecast horizons of 1, 4, and 7 months are summarized in Table 4. At the one-month horizon, the results are largely insignificant and the explained variation (R^2) is uniformly low, as expected, given the prior evidence of predictability only for longer horizons. Nonetheless, we note that about half of the signs for *VRP* are negative, while only the Swiss market produces a negative t-statistic for *NJRP*. Likewise, in the bivariate regressions, the t-statistic for *NJRP* is uniformly larger than the one for *VRP*. The latter feature, indicating superiority of the *NJRP* relative to *VRP*, is a robust feature observed across all indices and horizons.

At the four-month horizon, we start finding systematic indications of significance. The *VRP* remains a weak predictor, but the regression coefficient is positive for all indices except the Swiss. Most importantly, the *NJRP* regression coefficient is now significant in the majority of cases, and the reduction in explanatory power as we move from the bivariate regression to the univariate *NJRP* regression is very small, except for the Swiss index. Moreover, the R^2 values are moderately high at 7% – 12% for five of the indices.

	SP	ESTOXX	DAX	SMI	FTSE	MIB	IBEX							
Panel A: $R_{t,t+1}$														
VRP	-0.59	-1.08	0.88	-0.88	1.46	0.41	-2.58	-1.87	0.49	-0.04	-0.01	-1.77	0.63	-1.30
$VRP - NJRP$	-0.90	0.86	0.23	-0.49	1.50	1.32	-2.98	0.32	-0.87	0.04	0.80	0.79	2.35	1.64
$NJRP$	0.00	0.01	0.02	0.00	0.02	0.02	0.03	0.03	0.01	0.00	0.00	0.01	0.00	0.00
R^2														
Panel B: $R_{t,t+4}$														
VRP	0.45	-2.02	0.86	-3.22	2.18	-0.47	-3.35	-6.44	1.03	-0.06	0.31	-1.40	0.36	-2.34
$VRP - NJRP$	-0.03	2.49	1.16	-1.84	2.84	2.14	-7.34	2.23	-0.22	0.08	2.09	1.96	2.87	2.18
$NJRP$	0.00	0.00	0.04	0.03	0.01	0.11	0.09	0.02	0.00	0.01	0.00	0.07	0.00	0.01
R^2														
Panel C: $R_{t,t+7}$														
VRP	2.55	-0.64	2.03	-0.72	3.66	1.10	-0.57	-2.97	2.51	0.82	2.14	0.07	1.16	-1.46
$VRP - NJRP$	2.31	2.44	2.47	-0.15	3.67	2.87	-1.73	2.02	1.08	1.09	3.50	3.74	1.95	0.43
$NJRP$	0.06	0.03	0.11	0.11	0.05	0.00	0.17	0.16	0.05	0.01	0.03	0.01	0.06	0.00
R^2														

Table 4: **Predictive Regression.** For each index, the first column reports the t-statistic and the R^2 from the univariate regression of future returns on VRP , while the following three columns report the identical statistics for related regressions of future returns on risk premiums. In the second column, the single regressor is the variance risk premium minus the risk premium associated with negative jumps, the third column refers to a bivariate setting featuring both VRP and $NJRP$ as regressors, and in the fourth column the single regressor is the negative jump risk premium. Panel A displays the results for the future one-month returns, Panel B for four-month returns, and Panel C future seven-month returns.

The above tendencies only strengthen at the seven-month horizon. Except for the Swiss index, the R^2 rise considerably, and the majority of the $NJRP$ coefficients are significant. In addition, the exclusion of VRP from the bivariate regression has a negligible impact in terms of explanatory power. This is consistent with the extremely low R^2 statistics for the univariate $VRP - NJRP$ regression, where the negative jump risk premium is stripped from the variance risk premium. This effectively annihilates the predictive power associated with the VRP measure.

In summary, the majority, if not all, forecast power for the future equity risk premium resides with the negative jump risk premium, while the variance risk premium has negligible explanatory power, once the contribution from risk pricing of the left tail is netted out. These conclusions mirror the findings regarding the relative predictive power of the tail factor NJV^\perp and volatility factor V in Figure 5, although they appear somewhat less significant and consistent in Table 4.¹⁸

We also note that our predicted rank-ordering of explanatory power vis-a-vis the equity premium is confirmed. Comparing Table 4 and Figure 5, we find that NJV^\perp uniformly performs better at the 4-month horizon than the $NJRP$. At one month, the results are similar, but both measures have negligible predictive power. Finally, for the longer 7-month horizon, where the reliability of the inference may be questionable, the two predictors perform essentially on par. In all cases, the VRP produces substantially worse forecasts. Finally, we verify, see Section A.9 of the Appendix, that the VIX index has trivial explanatory power and only sporadically turns significant at the longest horizons. Although the associated R^2 -statistics do exceed those for the V factor, both variables are ineffective predictors of the equity risk premium.

Overall, the qualitative results are remarkably similar across indices, in spite of the large discrepancies in cumulative returns they experience, as highlighted in Figure 1. In other words, the predictive association between option-implied factors and the future realized returns and return variation appears to be operative almost uniformly for the major U.S. and European equity-index and derivatives markets. Moreover, Figures 1–3 suggest a substantial degree of coherence across the markets, although interesting country-specific features also are apparent. Hence, generally speak-

¹⁸The particularly poor forecast performance for the Swiss SMI index can be explained, in part, by the imposition of a cap on the Swiss franc-euro exchange rate on September 6, 2011, which induced an unexpected large one-time upward jump in the equity index. This observation constitutes a highly influential observation for the regression coefficients governing the forecasts for the actual future return variation. This suggests that the option-implied factors should possess superior predictive power relative to the risk premiums due to the additional noise associated with the forecasts for the future return variation. This is consistent with the discrepancy between the results for the SMI index in Figure 5 and Table 4.

ing, while some indices are subject to larger and more frequent negative shocks than others during the sample period, there is a striking similarity in the response to such shocks across markets. In particular, they have roughly analogous consequences for risk pricing, as manifest in their impact on option valuation across the moneyness-maturity spectrum, and consequently also on their extracted factor realizations. Through this channel, they induce a fairly high degree of correlation in signals for the direction of the future index appreciation and the associated return variation.

7 On Commonality in Risk and Risk Pricing

The observed correlation among the extracted option factors suggest that the markets are well integrated and the risk exposures quite similar. Thus, it is natural to consider a structure involving a combination of global and country-specific factors. However, for this to be feasible, we must take a stand on the nature of currency risk and its relation with equity risk. Obviously, the simplest approach is to ignore currency risk, but this is at odds with stylized facts regarding the reaction of global markets to tail events. For example, it is well-known that specific currencies take on “safe haven” characteristics during crises, e.g., the Swiss franc and partially the U.S. dollar. This feature implies that there may be a substantial degree of systematic and time-varying pricing of global risk embodied in the relative currency valuations.¹⁹

A different approach is to follow the commonly adopted assumption of effectively complete markets and determine the foreign exchange rate returns as the gap between the log-pricing kernels of the two countries. However, note that we purposely avoid providing a full representation of the pricing kernels, as we only specify the risk-neutral, and not the corresponding physical, return dynamics. As argued earlier, this approach ensures robustness against misspecification of the pricing kernel, i.e., we let the data speak more freely regarding the risk premiums. In addition, of course, this representation of the foreign exchange returns is invalid in the case of market incompleteness, see, e.g., Bakshi et al. (2017). This is particularly important in our case, where the focus is on rare

¹⁹Another concern is the lack of synchronous observations for the options and futures. The opening hours for the relevant exchanges vary widely, as reported in Table A1 of the Appendix. Since our sample contains tumultuous episodes associated with the financial crisis and the European sovereign debt crises, some influential return and option observations are recorded during trading hours on one exchange, but belong to the overnight period for other exchanges. In particular, during the sovereign debt crises, developments in the early trading hours on European markets, impacting the realized variation measures obtained from the futures markets, will be absorbed into the overnight return for the U.S. market, and vice versa for events in the late trading hours of the American market during the financial crisis. This will further hamper, in particular, the comparison of realized and expected future return variation measures, with direct consequences also for the associated estimates of the risk premiums.

events and their pricing.²⁰

With these caveats in mind, we limit ourselves to a simple correlation analysis that should help elicit any commonalities in the expected statistical and risk-neutral return variation measures across indices, and whether these relations remain stable across time. Table 5 reports, for two subsamples, the pairwise correlation of the expected negative jump variation and expected continuous return variation, where the latter is approximated by the total return variation net of the expected negative jump variation. These objective, or statistical, expected return variation measures display a high degree of correlation. The most notable change is a reduction in the correlation of the Spanish return variation measures with the remaining indices during 2010-2014 relative to 2007-2009. Beyond this feature, we note a relatively weak coherence between the *NJV* of Italy and the other indices in the initial subsample. In sum, the individual indices, denominated in their own currencies, appear to confront quite similar objective jump and volatility risk perceptions, with some indications that the Italian and Spanish indices display different risk exposures over the sample.

Table 6 speaks to the degree of commonality in the pricing of risk by reporting pairwise correlations of the estimated premiums associated with the one-month left jump tail and continuous return variation, respectively. The top panel shows that the compensation for negative jump risk is highly correlated across the indices in both sub-periods, with only the Spanish index deviating moderately from the remainder in the second period. The latter reflects the elevated downside risk pricing in Spain during the sovereign debt crises. In contrast, the correlation of the premium associated with the continuous return variation, proxied by the total variance risk premium minus the contribution of the negative jump risk premium, changes dramatically across the two subsamples. Initially, during 2007-2009, the compensation for (diffusive) variance risk is near identical across the indices, as reflected in the very high correlations reported in the left side of the lower panel of Table 6. In contrast, many of the pairwise correlations are almost negligible in the subsequent period. Hence, as the sovereign debt crisis unfolds, the option surfaces start to display distinct country-specific features, reflecting much stronger heterogeneity in the pricing of (continuous) variance risk.

Given the divergence in the broader risk pricing over the last subsample, the coherence in the left tail risk pricing is remarkable. It suggests a strong degree of commonality in the attitude towards downside equity risk, generating a strong correlation in equity market performance. In fact,

²⁰We further note that if so-called quanto options (U.S. dollar-denominated options on foreign equity indices) were available, we could circumvent these issues and directly extend our approach to a truly multivariate model.

$NJV_{t,t+30}^{\mathbb{P}}$												
2007-2009						2010-2014						
	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE
SP500		0.96	0.91	0.85	0.95	0.90	0.84		0.94	0.88	0.72	0.95
ESTOXX			0.98	0.92	0.99	0.86	0.93			0.95	0.80	0.97
DAX				0.95	0.98	0.78	0.96				0.88	0.95
SMI					0.94	0.66	0.97					0.85
FTSE						0.81	0.94					0.81
MIB							0.69					0.78
IBEX												0.72

$QV_{t,t+30}^{\mathbb{P}} - NJV_{t,t+30}^{\mathbb{P}}$												
2007-2009						2010-2014						
	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE
SP500		0.98	0.97	0.96	0.99	0.92	0.93		0.96	0.90	0.88	0.97
ESTOXX			0.99	0.96	0.99	0.95	0.97			0.95	0.90	0.98
DAX				0.95	0.98	0.93	0.96				0.96	0.94
SMI					0.97	0.91	0.95					0.92
FTSE						0.93	0.96					0.87
MIB							0.89					0.77
IBEX												0.85

Table 5: NJV and $RV - NJV$ Correlation

$NJRP_{t,t+30}$											
2007-2009						2010-2014					
SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE
	0.94	0.73	0.98	0.77	0.95	0.94		0.91	0.62	0.84	0.83
ESTOXX		0.90	0.94	0.91	0.97	0.97			0.85	0.90	0.95
DAX			0.75	0.95	0.82	0.86				0.84	0.88
SMI				0.78	0.92	0.94					0.88
FTSE					0.82	0.86					0.86
MIB						0.95					
IBEX											0.89
$VRP_{t,t+30} - NJRP_{t,t+30}$											
2007-2009						2010-2014					
SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE
	0.94	0.90	0.93	0.70	0.96	0.90		0.19	0.19	0.24	0.46
ESTOXX		0.94	0.98	0.82	0.97	0.97			0.66	-0.12	0.30
DAX			0.92	0.82	0.94	0.93				0.24	0.44
SMI				0.78	0.97	0.99					0.42
FTSE					0.72	0.76					0.40
MIB						0.96					
IBEX											0.80

Table 6: $NJRP$ and $VRP - NJRP$ Correlation

it is evident from Figure 1 that the return correlation is very strong, i.e., the periods of market appreciation and depreciation are highly synchronized. Thus, despite the dramatic differences in overall return performance, judging by the correlation across the option tail factors and the ensuing stock market returns, we assert that the ex-ante pricing of equity risk was comparable across the indices throughout the sample period. As such, these markets appear well integrated in terms of the required level of compensation for systematic risk.

8 Conclusion

This paper applies the option pricing approach of Andersen et al. (2015a) to a number of international equity indices, including the U.S. and various European derivatives markets. For all indices, there is a clean separation between a left tail factor, with predictive power for the future equity risk premiums, and a spot variance factor which is a potent predictor of the actual future return variation, but without explanatory power for future equity returns. Standard approaches, exploiting only volatility factors, miss the equity risk premium information in the option surface insofar as the volatility does not span the “pure tail factor,” which is the one embedding the predictive content for the equity risk premium.

We further document that the variance risk premium only has predictive power for the future equity returns due to the inclusion of the negative jump risk premium within the measure. Once the compensation for jump risk is stripped out of the variance risk premium, the explanatory power for the equity risk premium vanishes, implying that the signal concerning future returns stems primarily, if not exclusively, from the option-implied left jump factor. We also document that the left tail factor and the negative jump risk premium remain highly correlated across indices in the period following the financial crisis and through the European sovereign debt crises. In contrast, the volatility factor and diffusive volatility risk premium display a sharp drop in correlation across indices in the second part of the sample. This suggests a strong degree of commonality in the pricing of equity risk internationally, linked to the relative strength of the risk-neutral left jump intensity, whereas the underlying market risks at times vary markedly.

Finally, the characteristics of risk pricing in Italy and Spain deviate from the other indices over extended periods of time. It will be of interest, in future work, to relate the inferred downside tail risk premium in these countries to the economic events affecting them throughout this period.

Corresponding studies of the relative risk pricing across equity indices denominated in different currencies will require explicit consideration of currency risk. More generally, it will be useful to integrate the pricing of currency risk with our international stock market analysis.

References

- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and Forecasting Realized Volatility. *Econometrica* 71, 579–625.
- Andersen, T. G., O. Bondarenko, and M. T. Gonzalez-Perez (2015). Exploring Return Dynamics via Corridor Implied Volatility. *Review of Financial Studies* 28, 2902–2945.
- Andersen, T. G., N. Fusari, and V. Todorov (2015a). Parametric Inference and Dynamic State Recovery from Option Panels. *Econometrica* 83, 1081–1145.
- Andersen, T. G., N. Fusari, and V. Todorov (2015b). The Risk Premia Embedded in Index Options. *Journal of Financial Economics* 117, 558–584.
- Andersen, T. G., N. Fusari, and V. Todorov (2017). Pricing Short-Term Market Risk: Evidence from Weekly Options. *Journal of Finance* 72, 1335–1386.
- Bakshi, G., C. M., and C. J. (2017). Implications of Incomplete Markets for International Economies. *Review of Financial Studies*, forthcoming.
- Bakshi, G. and D. Madan (2000). Spanning and Derivative Security Valuation. *Journal of Financial Economics* 55, 205–238.
- Bakshi, G., G. Panayotov, and G. Skoulakis (2011). Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios. *Journal of Financial Economics* 100, 475–495.
- Bates, D. S. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies* 9, 69–107.
- Bates, D. S. (2000). Post-’87 Crash Fears in S&P 500 Future Options. *Journal of Econometrics* 94, 181–238.
- Bates, D. S. (2003). Empirical Option Pricing: A Retrospection. *Journal of Econometrics* 116, 387–404.
- Bekaert, G. and R. Hodrick (1992). Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets. *Journal of Finance* 47, 467–509.

- Bollerslev, T., J. Marrone, L. Xu, and H. Zhou (2014). Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence. *Journal of Financial and Quantitative Analysis* 49, 633–661.
- 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 (2011). Tails, Fears and Risk Premia. *Journal of Finance* 66, 2165–2211.
- Bollerslev, T., V. Todorov, and L. Xu (2015). Tail Risk Premia and Return Predictability. *Journal of Financial Economics* 118, 113–134.
- Broadie, M., M. Chernov, and M. Johannes (2007). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance* 62, 1453–1490.
- Campbell, J. and Y. Hamao (1992). Predicting Stock Returns in the United States and Japan: A Study of Long-Term Capital Market Integration. *Journal of Financial Markets* 47, 43–69.
- Carr, P. and L. Wu (2009). Variance Risk Premia. *Review of Financial Studies* 22, 1311–1341.
- Christoffersen, P., K. Jacobs, and K. Mimouni (2010). Models for S&P 500 Volatility Dynamics: Evidence from Realized Volatility, Daily Returns, and Option Returns. *Review of Financial Studies* 23, 3141–3189.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai (2012). Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options. *Journal of Financial Economics* 106, 447–472.
- Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics* 7, 174–196.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.
- Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Ferson, W. and C. Harvey (1993). The Risk and Predictability of International Equity Returns. *Review of Financial Studies* 6, 527–566.
- French, K., W. Schwert, and R. Stambaugh (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics* 19, 3–29.

- 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.
- Gruber, P., C. Tebaldi, and F. Trojani (2015). The Price of the Smile and Variance Risk Premia. Working paper, University of Lugano and University of Geneva.
- Harvey, C. (1991). The World Price of Covariance Risk. *Journal of Finance* 46, 111–157.
- Hjalmarsson, E. (2010). Predicting Global Stock Returns. *Journal of Financial and Quantitative Analysis* 45, 49–80.
- Jones, C. (2003). The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets. *Journal of Econometrics* 116, 181–224.
- Kou, S. (2002). A Jump- Diffusion Model for Option Pricing. *Management Science* 48, 1086–1101.
- Li, J. and G. Zinna (2017). The Variance Risk Premium: Components, Term Structures, and Stock Return Predictability. *Journal of Business Economics and Statistics*, forthcoming.
- Londono, J. and H. Zhou (2017). Variance Risk Premiums and the Forward Premium Puzzle. *Journal of Financial Economics* 124, 415–440.
- Merton, R. (1976). Option Pricing when Underlying Asset Returns are Discontinuous. *Journal of Financial Economics* 3, 125–144.
- Mueller, P., P. Sabtchevsky, A. Vedolin, and P. Whelan (2017). Variance Risk Premia on Stocks and Bonds. Working paper.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Phillips, P. (2014). On Confidence Intervals for Autoregressive Roots and Predictive Regression. *Econometrica* 82, 1177–1195.
- Santa-Clara, P. and S. Yan (2010). Crashes, Volatility and the Equity Premium: Lessons from S&P 500 Options. *Review of Economics and Statistics* 92, 435–451.
- Sizova, N. (2016). Do Long-Run Relations Convey Extra Information? Working paper, Rice University.
- Stambaugh, R. (1999). Predictive Regressions. *Journal of Financial Economics* 54, 783–820.
- Wang, H., H. Zhou, and Y. Zhou (2013). Credit Default Swap Spreads and Variance Risk Premia. *Journal of Banking and Finance* 37, 3733–3746.

A Appendix

A.1 Options Data

Index Name	Country	Exchange	OptionMetrics	Trading Hours	Tick Size	Multiplier
North America						
SP	USA	CBOE	SPX	8:30 am - 3:15 pm	0.05	100 \$
Europe						
ESTOXX	Europe	EUREX	SX5E	8:50 am - 5:30 pm	0.1	10 €
DAX	Germany	EUREX	DAX	8:50 am - 5:30 pm	0.1	5 €
SMI	Switzerland	EUREX	SMI	8:50 am - 5:20 pm	0.1	10 CHF
FTSE	UK	EURONEXT	UKX	8:00 am - 4:30 pm	0.5	10 £
MIB	Italy	IDEM	MIB	9:00 am - 5:40 pm	1.0	1 €
IBEX	Spain	MEFF	IBEX	9:00 am - 5:35 pm	1.0	2.5 €

Table A1: **Option Contract Specifications.** For each option contract we report the underlying index, the corresponding country, the name of the exchange, the symbol in the OptionMetrics database, and finally the trading hours, the tick size and the multiplier (as of December 2014).

A.2 Equity-Index Futures Data

Index Name	Country	Exchange	TickData	Daily Trading Hours	Tick Size	Multiplier
North America						
SP	USA	CME	ES	8:30 a.m. - 3:15 p.m.	0.25	50 \$
Europe						
ESTOXX	Europe	EUREX	XX	8:00 a.m. -10:00 p.m.	1.0	10 €
DAX	Germany	EUREX	DA	8:00 a.m. -10:00 p.m.	0.5	25 €
SMI	Switzerland	EUREX	SW	8:00 a.m. -10:00 p.m.	1.0	10 CHF
FTSE	UK	EURONEXT	FT	8:00 a.m. - 9:00 p.m.	0.5	10 £
MIB	Italy	IDEM	II	9:00 a.m. - 5:40 p.m.	0.5	10 £
IBEX	Spain	MEFF	IB	9:00 a.m. - 8:00 p.m.	0.5	10 £

Table A2: **Futures Contract Specifications.** For each index futures contract we report the country, the option exchange, the TickData symbol for the contract, the daily trading hours, the tick size, and the multiplier (as of December 2014).

A.3 Construction of High-Frequency Measures

The high-frequency futures data is obtained from TickData via the TickWrite software. We select Time Based Bars interval with one-minute granularity holding the last value in case there is no price change over the one-minute interval. We use the front maturity futures contract and we use the Auto Roll method provided by the software to roll-over to the next maturity when the front-maturity contract is near to expiration. We only consider the daily trading hours, combining pit and electronic trading. We apply the following filters to clean the raw data:

- We keep only observations between Monday and Friday.
- We remove days with no price changes and days corresponding to US holidays.
- We remove days with number of observations less than the average number of daily observations in the sample. This filter removes half-trading days such as the day before Thanksgiving or the day before Christmas for the US market.

The cleaned futures data is aggregated to five-minute frequency. For the construction of the high-frequency measures we introduce the following auxiliary notation. A generic series observed at high-frequency is denoted with Z (log futures price in our case). For ease of exposition, we ignore the overnight periods and assume that we have equidistant observations on the grid $0, \frac{1}{n}, \frac{2}{n}, \dots$. We denote $\Delta_n = \frac{1}{n}$ and $\Delta_i^n Z = Z_{\frac{i}{n}} - Z_{\frac{i-1}{n}}$. With this notation, the construction of the high-frequency measures is done following these steps:

1. Realized Variation:

$$RV_{t,t+\tau} = \sum_{i=\lfloor nt \rfloor + 1}^{\lfloor n(t+\tau) \rfloor} (\Delta_i^n Z)^2,$$

and if the interval $[t, t+\tau]$ includes an overnight period, the squared overnight return is added to the summation.

2. Bipower Variation:

$$BV_{t,t+\tau} = \frac{\pi}{2} \sum_{i=\lfloor nt \rfloor + 2}^{\lfloor n(t+\tau) \rfloor} |\Delta_i^n Z| |\Delta_{i-1}^n Z|.$$

3. Jumps Detection:

- For each series, we compute, the so-called Time-of-Day (TOD) function as defined in Bollerslev and Todorov (2011). We recompute the TOD function each time the exchange changed the trading hours.
- At each point in time, starting from the second day in the sample, we compute the threshold level:

$$\theta_i = \frac{3}{\sqrt{h}} \sqrt{RV_{t-h,t} \wedge BV_{t-h,t}} \times \Delta_n^{0.49} \times TOD_{i-\lfloor i/n \rfloor},$$

where h stands for 24 hours that exclude the overnight period.

- The increment $\Delta_i^n Z$ is flagged as one with jump if $|\Delta_i^n Z| > \theta_i$.

4. Truncated Variation :

$$TV_{t,t+\tau} = \sum_{i=\lfloor nt \rfloor + 1}^{\lfloor n(t+\tau) \rfloor} (\Delta_i^n Z)^2 1_{\{|\Delta_i^n Z| \leq \theta_i\}}.$$

5. Positive and Negative Jump Variation:

$$PJV_{t,t+\tau} = \sum_{i=\lfloor nt \rfloor + 1}^{\lfloor n(t+\tau) \rfloor} (\Delta_i^n Z)^2 1_{\{\Delta_i^n Z > \theta_i\}}, \quad NJV_{t,t+\tau} = \sum_{i=\lfloor nt \rfloor + 1}^{\lfloor n(t+\tau) \rfloor} (\Delta_i^n Z)^2 1_{\{\Delta_i^n Z < -\theta_i\}}.$$

6. For the computation of the local continuous variation, \widehat{V}_t^n , used to penalize the option-based volatility estimate in the objective function in (6), we take into account the following:

- $\widehat{V}_t^n = \frac{1}{h} TV_{t,t+\tau}$, where τ is equal to 3 hours.
- If we encounter more than 4 consecutive zero returns (this could happen in case of market closure or "lunch break") then we extend the window until we reach 36 returns containing less than 4 consecutive zero returns.

To simplify the notation, RV_t , TV_t , PJV_t , and NJV_t refer to the above quantity computed on a daily basis (i.e. from the opening to the closing of the market).

We then use the HAR model of Corsi (2009) to forecast the future quadratic variation and

negative jump variation under the statistical measure:²¹

$$\begin{aligned}\ln(TV_{t+1,t+h}^{\mathbb{P}}) &= a + b_1 \log(TV_t) + b_2 \log(TV_{t-4,t-1}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(TV_{t-84,t-21}) + \epsilon_t \\ TV_{t,t+h}^{\mathbb{P}} &= \exp(a + b_1 \log(TV_t) + b_2 \log(TV_{t-4,t-1}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(TV_{t-84,t-21}) + \sigma_\epsilon^2/2),\end{aligned}\quad (12)$$

where σ_ϵ^2 is the variance of the residual from the linear regression in Equation (12). Equivalently, for the negative jump variation we have:

$$\begin{aligned}\ln(NJV_{t+1,t+h}^{\mathbb{P}}) &= a + b_1 \log(TV_t) + b_2 \log(TV_{t-4,t-1}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(TV_{t-84,t-21}) + \epsilon_t \\ NJV_{t,t+h}^{\mathbb{P}} &= \exp(a + b_1 \log(TV_t) + b_2 \log(TV_{t-4,t-1}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(TV_{t-84,t-21}) + \sigma_\epsilon^2/2),\end{aligned}\quad (13)$$

while for the positive jump variation we have

$$\begin{aligned}\ln(PJV_{t+1,t+h}^{\mathbb{P}}) &= a + b_1 \log(TV_{t-1}) + b_2 \log(TV_{t-4,t}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(PJV_{t-20,t-5}) + \epsilon_t \\ PJV_{t+1,t+h}^{\mathbb{P}} &= \exp(a + b_1 \log(TV_t) + b_2 \log(TV_{t-4,t-1}) + b_3 \log(TV_{t-20,t-5}) + b_4 \log(TV_{t-84,t-21}) + \sigma_\epsilon^2/2).\end{aligned}\quad (14)$$

Finally, the quadratic variation is obtained as:

$$QV_{t+1,t+h}^{\mathbb{P}} = TV_{t+1,t+h}^{\mathbb{P}} + NJV_{t+1,t+h}^{\mathbb{P}} + PJV_{t+1,t+h}^{\mathbb{P}}.$$

A.4 The Model-Free Option-Implied VIX Measure

The measure is computed following Carr and Wu (2009). On each day, we take the two option cross-sections with tenor closest to 30 calendar days. For each cross section, we create a fine grid of strike prices K covering the moneyness range defined as $-8 \leq m \leq 8$ with increments of 0.1. We then interpolate the implied volatility as a function of K . When K is lower (higher) than the lowest (highest) available strike, we extrapolate the implied volatility outside the defined moneyness range as a constant equal to the implied volatility at the lowest (highest) available strike. We obtain the VIX index for both maturities and linearly interpolate to obtain the VIX index corresponding to 30 calendar days.

A.5 Sensitivity of the Jump Factor to the Volatility Specification

As discussed earlier, our model is parsimonious and, in particular, contains only a single volatility factor. One potential concern is that the relatively constrained modeling of the volatility process may induce a bias in the extraction of the implied jump dynamics from the option surface. For the S&P 500 sample, we have a natural benchmark. Andersen et al. (2015b) estimate an extended version of model (1) which provides an excellent fit to the option prices as well as the volatility series, clearly outperforming more standard and equally heavily parameterized representations.

²¹When estimating the HAR model we appropriately rescale each realized measure in order to match the unconditional daily variance level. Specifically we multiply each quantity by $\overline{y_t^2}/\overline{RV}_t$, where $\overline{y_t^2}$ is the average squared daily return (close-to-close) and \overline{RV}_t is the average level of the realized variance computed during the trading time.

The sample exploited here is shorter than in Andersen et al. (2015b), due to the synchronization with the data available for the European markets. On the other hand, it covers a wider cross-section since we include the full set of short-dated (weekly) options, which were excluded from the analysis in Andersen et al. (2015b). Moreover, the current sample extends through December 2014, while the elaborate model is estimated over January 1996–July 21, 2010, but with out-of-sample extraction of the jump intensity factor through April 23, 2013. Hence, the two series of option-implied jump factors overlap over the period January 2007 till April 2013.

Because the left tail factor enters the specification for the jump intensity in distinct ways, a direct comparison of the factor realizations across the two models is not meaningful. Instead, we focus on the model-implied variation in the negative jump intensity stemming from the left tail factor. In both cases, the effect is proportional to the concurrent value of U factor in the model. Consequently, Figure A1 depicts the standardized jump factor series extracted from the two separate models, estimated from partially overlapping periods and option samples. The models deliver remarkably similar time series paths for the return variation induced by the left tail factor, with a correlation of 0.987 during the overlap, corroborating the robust identification of the jump tail factor from either model.




Figure A1: Comparison of Model-Implied Left Jump Tail Factors. The figure depicts the daily standardized option-implied left jump tail factors from model (1) and the Andersen et al. (2015b) model. The series are extracted based on parameter estimates from weekly SPX option prices observed over January 2007–December 2014 and January 1996–July 21, 2010, respectively, with the latter series having been extended in Andersen et al. (2015b) to cover jump intensity factors for the out-of-sample period, July 22, 2010 - April 23, 2013, as well. Both series are normalized to have mean zero and unit variance.

A.6 Parameter Estimates

A.6.1 S&P500

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.012	c_0^+	7.551	0.622
\bar{v}	0.029	0.000	λ^-	13.548	0.196
κ	6.216	0.086	λ^+	78.255	2.368
σ	0.605	0.008	μ_v	22.663	0.462
κ_u	0.600	0.062	μ_u	18.656	5.389
Panel B: Summary Statistics					
RMSE				1.909	
Mean negative jump intensity				0.853	
Mean negative jump size				-0.074	
Mean positive jump size				0.013	
Mean diffusive variance				0.038	
Mean negative jump variance				0.009	
Mean positive jump variance				0.002	

Table A3: **Estimation Results for the Parametric Model (1) Fit to the S&P 500 Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.2 ESTOXX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.999	0.014	c_0^+	12.879	0.777
\bar{v}	0.032	0.000	λ^-	14.569	0.180
κ	7.471	0.170	λ^+	66.123	1.423
σ	0.687	0.011	μ_v	13.089	0.263
κ_u	1.511	0.050	μ_u	154.224	8.638
Panel B: Summary Statistics					
RMSE				1.974	
Mean negative jump intensity				1.632	
Mean negative jump size				-0.069	
Mean positive jump size				0.015	
Mean diffusive variance				0.054	
Mean negative jump variance				0.015	
Mean positive jump variance				0.006	

Table A4: **Estimation Results for the Parametric Model (1) Fit to the ESTOXX Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.3 DAX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.012	c_0^+	9.666	0.471
\bar{v}	0.024	0.000	λ^-	18.967	0.189
κ	11.653	0.146	λ^+	68.833	1.148
σ	0.741	0.010	μ_v	32.642	0.690
κ_u	0.262	0.033	μ_u	44.814	6.474
Panel B: Summary Statistics					
RMSE				2.050	
Mean negative jump intensity				1.791	
Mean negative jump size				-0.053	
Mean positive jump size				0.015	
Mean diffusive variance				0.055	
Mean negative jump variance				0.010	
Mean positive jump variance				0.004	

Table A5: **Estimation Results for the Parametric Model (1) Fit to the DAX Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.4 SMI

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.752	0.010	c_0^+	4.397	0.445
\bar{v}	0.020	0.000	λ^-	24.316	0.425
κ	12.394	0.319	λ^+	66.762	2.024
σ	0.712	0.013	μ_v	27.147	0.579
κ_u	1.724	0.090	μ_u	415.332	41.169
Panel B: Summary Statistics					
RMSE				1.835	
Mean negative jump intensity				2.649	
Mean negative jump size				-0.041	
Mean positive jump size				0.015	
Mean diffusive variance				0.032	
Mean negative jump variance				0.009	
Mean positive jump variance				0.002	

Table A6: **Estimation Results for the Parametric Model (1) Fit to the SMI Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.5 FTSE

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.014	c_0^+	5.099	0.979
\bar{v}	0.028	0.000	λ^-	12.688	0.389
κ	5.381	0.138	λ^+	66.693	5.117
σ	0.546	0.008	μ_v	30.845	1.347
κ_u	1.191	0.117	μ_u	25.114	15.159
Panel B: Summary Statistics					
RMSE				2.164	
Mean negative jump intensity				0.604	
Mean negative jump size				-0.079	
Mean positive jump size				0.015	
Mean diffusive variance				0.040	
Mean negative jump variance				0.008	
Mean positive jump variance				0.002	

Table A7: **Estimation Results for the Parametric Model (1) Fit to the FTSE Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.6 MIB

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.973	0.038	c_0^+	1.564	0.250
\bar{v}	0.029	0.000	λ^-	10.823	0.284
κ	7.111	0.201	λ^+	24.902	1.458
σ	0.501	0.018	μ_v	24.804	1.026
κ_u	0.512	0.078	μ_u	20.066	6.613
Panel B: Summary Statistics					
RMSE				2.867	
Mean negative jump intensity				0.807	
Mean negative jump size				-0.092	
Mean positive jump size				0.040	
Mean diffusive variance				0.063	
Mean negative jump variance				0.014	
Mean positive jump variance				0.005	

Table A8: **Estimation Results for the Parametric Model (1) Fit to the MIB Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.6.7 IBEX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.009	c_0^+	1.500	0.350
\bar{v}	0.040	0.001	λ^-	17.339	1.104
κ	6.530	0.239	λ^+	23.695	2.251
σ	0.683	0.012	μ_v	6.147	1.835
κ_u	1.244	0.438	μ_u	180.477	119.791
Panel B: Summary Statistics					
RMSE				2.085	
Mean negative jump intensity				3.419	
Mean negative jump size				-0.058	
Mean positive jump size				0.042	
Mean diffusive variance				0.060	
Mean negative jump variance				0.023	
Mean positive jump variance				0.005	

Table A9: **Estimation Results for the Parametric Model (1) Fit to the IBEX Option Panel and High-Frequency Spot Volatility Measure.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or a neighboring business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. The RMSE is given in annualized BSIV terms, while the variances and jump intensity is given in annualized decimal units and, finally, the jump sizes are given in decimal units.

A.7 Predictive Regressions for the Negative Jump Activity

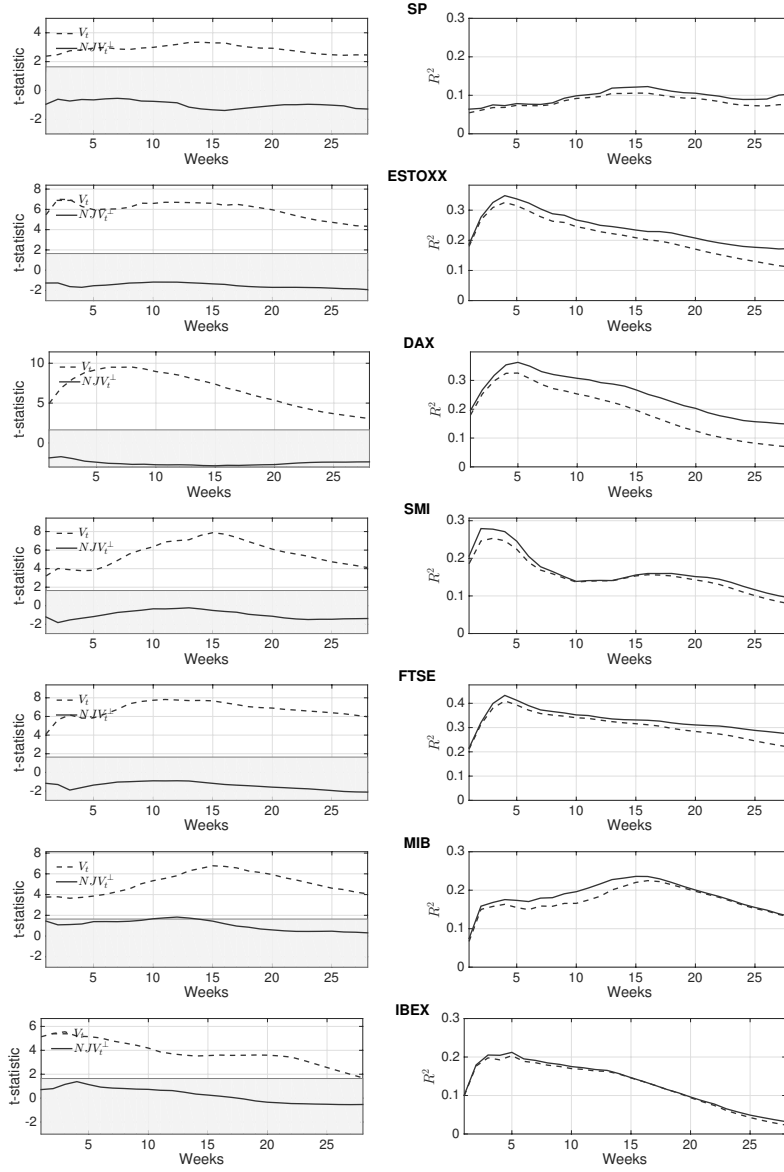


Figure A2: **Predictive Regressions for Negative Jumps.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line signifies the part explained by the spot variance alone.

A.8 Factor and Risk Premium Interactions

	V, NJV	VRP, V	VRP, NJV^\perp	$NJRP, V$	$NJRP, NJV^\perp$	$CVRP, V$	$CVRP, NJV^\perp$
2007-2009							
SP	0.86	0.90	0.32	0.79	0.72	0.89	0.01
ESTOXX	0.68	0.87	0.39	0.40	0.95	0.90	-0.08
DAX	0.70	0.80	0.37	0.07	0.99	0.92	-0.03
SMI	0.79	0.95	0.21	0.60	0.86	0.95	0.01
FTSE	0.80	0.64	-0.10	0.24	0.97	0.72	-0.38
MIB	0.82	0.93	0.05	0.72	0.79	0.89	-0.13
IBEX	0.49	0.93	0.09	0.35	0.85	0.94	-0.35
2010-2014							
SP	0.91	0.91	0.85	0.88	0.97	0.89	0.69
ESTOXX	0.84	0.58	0.77	0.71	0.98	0.57	0.19
DAX	0.83	0.14	0.57	0.23	0.92	0.60	0.30
SMI	0.77	0.30	0.61	0.21	0.97	0.63	-0.03
FTSE	0.83	0.51	0.45	0.58	0.97	0.51	0.02
MIB	0.80	0.84	0.44	0.68	0.90	0.81	0.13
IBEX	0.76	0.86	0.77	0.68	0.98	0.53	-0.24

Table A10: **Factor and Risk Premium Correlations.** The table reports the sample correlation of daily observations for a combination of option factors and risk premiums over two sub-periods, 2007-2009 and 2010-2014.




Figure A3: **Variance, Negative Jump Variation (orthogonal), and Variance Risk Premium.** Each panel shows the 21-day backward moving average of the option implied variance V_t (dotted line), the orthogonal part of the negative jump variation with respect to the option implied variance NJV_t^\perp (dashed line), and the variance jump risk premium $VRP_{t,t+30}$ (solid line). All series are standardized to have mean zero and unit variance.




Figure A4: **Variance, Negative Jump Variation (orthogonal), and Negative Jump Risk Premium.** The panels display the 21-day backward moving average of the option implied variance V_t (dotted line), the orthogonal part of the negative jump variation with respect to the option implied variance NJV_t^\perp (dashed line), and the negative jump risk premium $NJRP_{t,t+30}$ (solid line). All series are standardized to have mean zero and unit variance.

A.9 VIX and the Equity Risk Premium

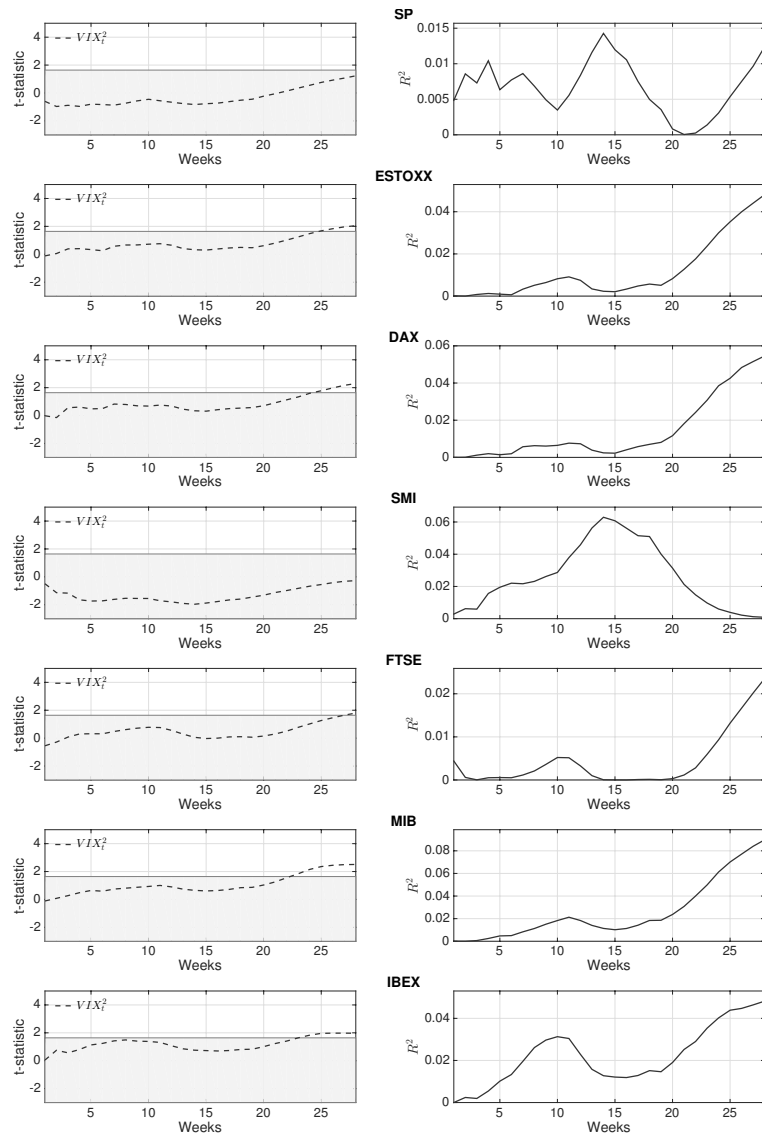


Figure A5: **VIX-Based Predictive Regressions for Excess Returns.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 .

Research Papers

2018



- 2017-25: Massimiliano Caporin, Gisle J. Natvik, Francesco Ravazzolo and Paolo Santucci de Magistris: The Bank-Sovereign Nexus: Evidence from a non-Bailout Episode
- 2017-26: Mikkel Bennedsen, Asger Lunde and Mikko S. Pakkanen: Decoupling the short- and long-term behavior of stochastic volatility
- 2017-27: Martin M. Andreasen, Jens H.E. Christensen and Simon Riddell: The TIPS Liquidity Premium
- 2017-28: Annastiina Silvennoinen and Timo Teräsvirta: Consistency and asymptotic normality of maximum likelihood estimators of a multiplicative time-varying smooth transition correlation GARCH model
- 2017-29: Cristina Amado, Annastiina Silvennoinen and Timo Teräsvirta: Modelling and forecasting WIG20 daily returns
- 2017-30: Kim Christensen, Ulrich Hounyo and Mark Podolskij: Is the diurnal pattern sufficient to explain the intraday variation in volatility? A nonparametric assessment
- 2017-31: Martin M. Andreasen, Jens H.E. Christensen and Glenn D. Rudebusch: Term Structure Analysis with Big Data
- 2017-32: Timo Teräsvirta: Nonlinear models in macroeconometrics
- 2017-33: Isabel Casas, Eva Ferreira and Susan Orbe: Time-varying coefficient estimation in SURE models. Application to portfolio management
- 2017-34: Hossein Asgharian, Charlotte Christiansen, Ai Jun Hou and Weining Wang: Long- and Short-Run Components of Factor Betas: Implications for Equity Pricing
- 2017-35: Juan Carlos Parra-Alvarez, Olaf Posch and Mu-Chun Wang: Identification and estimation of heterogeneous agent models: A likelihood approach
- 2017-36: Andrés González, Timo Teräsvirta, Dick van Dijk and Yukai Yang: Panel Smooth Transition Regression Models
- 2017-37: Søren Johansen and Morten Ørregaard Nielsen: Testing the CVAR in the fractional CVAR model
- 2017-38: Nektarios Aslanidis and Charlotte Christiansen: Flight to Safety from European Stock Markets
- 2017-39: Tommaso Proietti, Niels Haldrup and Oskar Knapik: Spikes and memory in (Nord Pool) electricity price spot prices
- 2018-01: Emilio Zanetti Chini: Forecaster's utility and forecasts coherence
- 2018-02: Torben G. Andersen, Nicola Fusari and Viktor Todorov: The Pricing of Tail Risk and the Equity Premium: Evidence from International Option Markets