

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

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

June 20, 2016

Abstract

The paper explores the global pricing of market tail risk as manifest in equity-index options. We document the presence of a left tail factor that displays large persistent shifts, largely unrelated to the corresponding dynamics of return volatility. This left tail factor is a potent predictor of future excess equity-index returns, while the implied volatility only forecasts future equity return variation, not the expected returns. We conclude that the option surface dynamics embed separate equity risk and risk premium factors which are successfully disentangled by our simple two-factor affine model for all the equity indices explored. The systematic deviations across countries speak to the differential risk and its pricing during the great recession and the European sovereign debt crises. Most strikingly, the Southern European indices are subject to much higher tail risk pricing following the onset of the sovereign debt crisis. Their stock markets react in a nearly identical fashion to Euro-wide systematic tail events. Nevertheless, these events are priced very differently by their option markets, suggesting difference in crash beliefs across the countries which are hard to detect with stock market data alone.

Keywords: Equity Risk Premium, Extreme Events, Jumps, Option Pricing, Return Predictability, Stochastic Volatility, Tails, Variance Risk Premium.

JEL classification: C51, C52, G12.

*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.

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

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

[§]Department of Finance, Kellogg School of Management, Northwestern University,1 Evanston, IL 60208; e-mail: 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 to accommodate the observed interdependencies between tail events and their pricing across markets.

In this paper, we argue that the increasing liquidity of derivative markets worldwide provide a unique opportunity to shed some light on this question. In particular, the trading in equity-index option markets has grown sharply in most 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 trade prices for financial securities that are highly informative 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 US (S&P 500), Euro-zone (ESTOXX), Germany (DAX), Switzerland (SMI), UK (FTSE), Italy (MIB), and Spain (IBEX) over 2007-2014 to extract factors that are critical in pricing risk across the main equity indices in North America and Europe.

The confluence of turbulent periods render 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 financial markets, and we exploit the options data to study how tail risk was perceived and priced across these episodes. At the same time, several of the European countries had unique exposures to the sovereign debt crises. It is of separate interest to determine whether these risks were priced similarly across the diverse markets. By combining the pricing of financial risks with ex post information on actual realized returns, return volatilities and jumps, we can gauge the size of the risk premiums and what factors drive the risk compensation. We notice striking heterogeneity in the stock market performance over the sample, with the German market appreciating by an average of close to 5% per year and the Italian index depreciating by 10% annually. We exploit this heterogeneity to strengthen earlier empirical evidence (Andersen et al. (2015b) and Bollerslev et al. (2015)) regarding the connection between market tail risk and the equity risk premium which pertains to the US market.

Standard option pricing models capture the dynamics of the equity-index option surfaces 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 by regular volatility factors. Hence, a distinct factor is necessary to account for the priced downside risk in the option surface relative to the regular

volatility factors. In empirical work focused on the US markets, this state variable has emerged as a crucial equity risk premium factor, see Andersen et al. (2015b).¹

Given the findings above, it would be natural to adopt the Andersen et al. (2015b) model for all indices explored in this study. The specification involves two volatility and a separate jump factor, thus encompassing most standard option pricing models in the literature, while also maintaining the left tail features that were found to be critical for the long sample of S&P 500 options. However, many of the European option samples are shorter and contain fewer options in the strike cross-section. As such, separate identification of the two volatility and the tail factor is challenging in practice. As a consequence, we pare down the model to a single volatility component and a tail factor. This facilitates robust identification of the factors and reduces the number of model parameters considerably, providing a solid basis for out-of-sample exploration of the predictive power of the factors. Further, we confirm that the tail factor extracted from our simplified set-up matches the one obtained from the original Andersen et al. (2015b) model very closely, while the volatility factor has high predictive power for the future return volatility and jump activity. In fact, 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 typically remains elevated long after market volatility subsides to pre-crisis levels. This feature is in line with the findings of Andersen et al. (2015b), and is not compatible with the usual approach to the modeling of volatility and jump risk in the option pricing literature.

This stark separation of the tail risk factor and market volatility has important implications for the pricing and dynamics of market risks. Although market volatility is a strong predictor of future market risks, 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. The tail factor also is 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, asset prices are heavily discounted and the option-based tail factor is elevated long after market volatility resides. This combination turns out to be a strong signal that the stock market will rebound in the future.

In terms of return volatility, we find the European and US 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 crisis. Overall, the UK, Swiss and, to some extent, the

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

German volatilities remain very close throughout the sample. The largest divergences in volatility dynamics occur between the US, UK and Swiss indices on one side and the Spanish and Italian ones 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 within the latter there is, at times, a stark discrepancy in the priced tail risk. Specifically, the Spanish tail factor reacts strongly to both major phases of the sovereign debt crisis, while the Italian response is more muted, especially for the first phase.

To separate the risk premium component of the tail risk factor, we complement the above option-based evidence with an analysis of high frequency return data. In particular, we compare the realized jump risks across the market indices. We focus attention in this analysis to the German, Italian and Spanish stock markets. Our analysis shows remarkable commonality in the way the three stock markets react to systematic jumps events (i.e., times at which all three markets jump simultaneously) over the period of 2007-2009, including the global financial crisis. On the other hand, in the second part of the sample, containing the European sovereign debt crises, we notice a more muted response of the German market index to systematic jump events than that of the Spanish and Italian markets. Meanwhile, the latter continue to have similar (i.e., statistically indistinguishable from one-to-one) reactions to systematic jump events during this period.

This analysis shows that the differences in the tail factors in Spain and Italy is not due to the prevalence of more frequent large negative return shocks in Spain. Instead, it seems as if the Spanish market was gripped by “fears” and oversized risk premiums, unlike those even in Italy. A potential explanation is the concern for a possible rupture of the eurozone area, with a subsequent creation of a “Southern euro” that may include Spain, but not Italy. This suggests the divergence in the option-based tail factors of the two countries reflects a so-called “peso problem” in the sense that the stock market return data is not sufficient to capture the crash beliefs of the market participants.

In summary, we find a striking degree of robustness in the pricing of financial market risks across the US and the European indices. Moreover, the compensation for such risks appear to evolve similarly across the countries, with the exception of the two Southern European nations. The result confirms that the option surfaces allow for separate identification of risk and risk premium factors, represented by the spot variance and left tail factors respectively. For all indices, we verify the predictive power of the tail factor for the future equity risk premium. This factor is also the primary determinant of the left tail risk premium and an important component of the variance risk premium. In contrast, only the spot variance factor provides predictive power for the actual future

return variation and jump activity. Thus, the separation into risk premium (left tail) and risk (volatility) factors is robust and economically informative. Our main conclusions elude standard option pricing models, where the requisite separation between the tail and volatility factors is lacking and the associated risk premiums are intertwined, and thus not suitably identifiable.

The rest of the paper is organized as follows. We start our analysis with summarizing our key empirical results for the US market index in Section 2. We review the data used in the paper in Section 3. Section 4 presents the model we use to fit the option surfaces and extract information from them, and in Section 5 we review the estimation method. Section 6 depicts the empirically extracted tail factors for all the indices. Section 7 analyzes the implications for the predictability of future risk premiums and return variability on international stock markets. In Section 8, we use the option-extracted factors and the high-frequency futures data to explore the premia for negative jump risk in the euro-zone countries. Section 9 concludes. Additional details on various aspects related to the data, the estimation and the results are contained in Section 10.

2 Illustrative Evidence for the S&P 500 Index

This section exemplifies some of our key findings by reviewing results obtained from S&P 500 equity-index (SPX) options at the Chicago Board Options Exchange (CBOE) and high-frequency returns on the e-mini S&P 500 futures at the CME Group. It extends evidence from Andersen et al. (2015b), but relies on a simplified two-factor parametric model for the risk-neutral distribution. Our sample is shorter, but contains a wider array of observations, as we include very short-dated options. As such, the results speak to the robustness of prior evidence and serve as a benchmark for the European indices explored below.

Our sample covers January 2007–December 2014 – the period for which we have a broad set of option quotes for the European countries. The risk-neutral return dynamics is governed by a traditional volatility factor and a jump factor with different intensities and amplitudes for positive and negative jumps. Finally, the negative return and positive volatility jumps coincide. Even if the model is significantly pared down relative to Andersen et al. (2015b), it captures the salient features of the option surface and equity-index return dynamics across time and for all the individual markets. Details regarding data sources, the parametric models and the inference techniques are deferred to subsequent sections.

2.1 The U.S. Evidence

We document robust features of the fluctuations in the option surface over time that speak to critical aspects of the pricing of the distinct factors governing the equity return dynamics.

One primary finding of Andersen et al. (2015b) is that the downside jump factor, as manifest in the option surface, does not obey a tight relationship with the stochastic volatility driving the underlying equity-index return dynamics. Specifically, the magnitude and fluctuations in the option-implied volatility skew cannot be captured adequately through traditional option pricing models where the dynamics is governed exclusively by volatility components. Instead, we exploit a tractable affine model with a separate downside jump intensity factor. For parsimony and ease of identification, our representation has a single volatility factor, V_t , and a separate downside jump intensity factor, U_t , driving the risk-neutral return dynamics. This provides a reasonable fit to the salient features of the option surface dynamics, it ensures strong identification of the two factors, and it facilitates the economic interpretation of the results.

Using our parametric model for the risk-neutral dynamics, we extract the implied realization of the volatility and jump factors day-by-day from the S&P 500 options over the full sample. The top panel of Figure 1 displays the end-of-trading-day implied (annualized) spot volatility. The extreme spike in volatility surrounding the financial crisis stands out, while the elevation associated with the two main stages of the European sovereign debt crisis starting in spring 2010 and summer 2011, respectively, are evident, but decidedly more muted.

The middle panel of Figure 1 depicts the corresponding negative risk-neutral jump variation, obtained as the (conditionally) expected number of negative jumps times the expected jump amplitude.² In contrast to spot volatility, the reaction of this tail measure to the European debt crisis is not dwarfed by the response to the 2008-09 financial crisis. Thus, while both series are highly elevated during crises, the magnitude of their response differs importantly across events.

In the bottom panel of Figure 1, we display the residual from the regression of the negative jump variation on the spot variance which we denote \tilde{U} . It captures the component of the negative jump variation that is unspanned by, i.e., not linearly related to, the spot variance. This further illustrates the differential response of the volatility and left jump factors across distinct turbulent episodes. At the initial phase of the financial crisis, volatility sky-rockets compared to the jump variation, inducing a huge negative inlier in the jump residual series. Apparently, the immediate reaction to the market crash in the option market was one of profound uncertainty rather than a predominant perception, or fear, of future abrupt negative return shocks. However, over the

²For compatibility with spot volatility, we plot the square-root of the annualized jump variation.




Figure 1: **Implied Spot Volatility and Negative Jump Volatility.** The figure displays the five-day trailing moving average of the spot volatility (top panel), the (square-root of the) negative jump variation (middle panel), and the residual of the negative jump variation regressed on the spot variance (bottom panel) implied by the E-mini S&P 500 futures options for 2007-2014. All measures are reported in annualized decimal units.

subsequent month the implied negative jump variation rises sharply and remains highly elevated even as spot volatility recedes. Nonetheless, the increase in the residual jump intensity only slightly exceeds what is experienced in the following European crisis, implying a substantially lower tail pricing relative to volatility in the former case.

The bottom panel also reveals a common trait emerging from these disruptive episodes. Following each, there is a prolonged period in which the left jump factor remains elevated relative to spot volatility. In other words, the “excitement” of the left tail lingers for much longer than for volatility. This phenomenon generates large disproportional shifts in the left part of the implied volatility surface, associated with rich pricing of out-of-the-money puts, relative to the level observed for other options.

In our model, two state variables – the spot variance and negative jump variation – govern the risk-neutral return dynamics. To explore the implications for risk premia, we must relate the state variables to the underlying conditional equity-index return distribution and, in particular, the expected excess returns and risk. In affine models, this relation is linear. We first consider the equity risk premium relation. To ensure genuine out-of-sample forecasting, we estimate the parameters of the option pricing model over the year 2007 alone.³ Conditional on these parameter estimates, the state variables are extracted day-by-day throughout the sample. Importantly, this implies that the state variables for the remainder of the sample, 2008-2014, are obtained from prior information only. Hence, when the state variables serve as explanatory variables for the future excess index returns over this period, there is no mechanical look-ahead bias.

Figure 2 depicts the significance of the regression coefficients and the degree of explanatory power for the state variables, i.e., the implied spot variance, V_t , and the residual jump intensity, \tilde{U}_t , for the future excess returns on the S&P 500 index in a simple bivariate predictive regression. Given our limited sample, we consider only horizons spanning one week to about half a year.

We find the excess returns to be linked to the component of the left jump tail intensity factor unspanned by V_t , i.e., \tilde{U}_t , for horizons beyond two months, with a maximum attained around four months. The explained variation in the excess returns exceeds 20% for horizons beyond 3 1/2 months. At the same time, the implied spot volatility is insignificant, so the information regarding the pricing of equity return risk is embedded solely within the negative jump factor.

While the results displayed in Figure 2 are striking, they are perhaps not entirely surprising given the evidence in Bollerslev et al. (2009) showing that the variance risk premium (VRP), and not the (spot) variance, predicts future equity index returns. The variance risk premium consists of

³Our estimation procedure is sufficiently efficient and the available option cross-section large enough to produce reasonably accurate point estimates, even for a limited time span.




Figure 2: **Predictive Regression for Excess Returns using the State Variables.** The figure reports findings from regressing future weekly excess equity-index returns on the option-implied state variables – the spot variance and negative jump intensity (orthogonal to spot variance). The left panel displays t-statistics for the regression coefficients. The right panel depicts the R^2 obtained including both explanatory variables (solid line) and excluding the jump factor (dashed line).

the option-implied return variance as, e.g., captured by the VIX index minus the expected future realized return variation. Our left tail factor is a component of the return variation measure that constitutes the VIX, so there is potentially a mechanical correlation between the two variables. To gauge whether the variance risk premium subsumes the explanatory power in the tail factor, we consider a predictive regression for the equity-index returns where the explanatory variables are the tail factor, \tilde{U}_t , and the component of the VRP that is orthogonal to the tail factor.




Figure 3: **Predictive Regression for Excess Returns using the Variance Risk Premium.** The figure reports findings from regressing future weekly excess equity-index returns on the option-implied negative jump intensity (orthogonal to spot variance) and variance risk premium (VRP). The VRP is estimated by non-parametric means as the difference between the CBOE VIX squared and the predicted monthly realized variance (plus overnight squared returns). The forecast for realized variance is obtained from the model of Corsi (2009). The left panel displays the t-statistic for the regression coefficients on the explanatory variables (\tilde{U} and VRP). The right panel depicts the R^2 obtained using just the VRP (dashed line) and both \tilde{U} and VRP (solid line).

Figure 3 confirms that the component of the variance risk premium unrelated to our tail factor may possess some predictive power for the equity risk premium at horizons beyond four months. However, the significance of the tail factor is much more pronounced, and the auxiliary contribution from the VRP to the overall explained variation is marginal. In other words, the left tail factor

emerges as the critical component embedded within the option surface for explaining the size of the (conditional) equity risk premium.

In contrast, when we regress the future (realized) return variation, stemming from both diffusive volatility and jumps, on the state variables, the relative explanatory power is reversed. The jump intensity unspanned by the market volatility has no predictive power for the actual future return variation, whereas the spot volatility is a strong predictor of future return volatility and jump variation. In short, the future return risks are well accounted for by the current volatility, which is identifiable from both the option surface and return observations on the underlying asset. But, as we have seen, this factor is unrelated to the equity risk premium which, instead, is tied to the part of the left jump tail factor orthogonal to volatility. These findings suggest a stark separation between equity market risk – as reflected by the expected future volatility and jumps – and the pricing of equity risk – as manifest in the average future market excess returns. Since the option tail factor appears detached from the actual return dynamics, it is infeasible to extract information regarding this factor from the return observations alone. That is, our results suggest that the option surface embodies critical information for the identification of the equity risk premium, corroborating the evidence in Andersen et al. (2015b).




Figure 4: **Predictive Regression for Return Variation.** The figure reports findings from regressions of future cumulative return variation on current option-implied state variables, i.e., the spot variance and negative jump intensity. The left panel displays t-statistics for the regression coefficients on the state variables. The right panel depicts the corresponding regression R^2 .

These findings raise a number of questions. Are similar option pricing and risk premium patterns present in other developed economies? Are the fundamental U.S. state variables related to corresponding factors in other countries? In other words, how universal are our results, and how may key features of the risk-neutral distributions be linked? We explore these topics later, but we first introduce our data sources and introduce our risk-neutral two-factor model more formally.

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 data for the underlying equity indices.

3.1 Equity-Index Option Data

Our option data are obtained from the new OptionMetrics Ivy DB Global Indices database that 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 currency are also provided. To reduce the computational burden, we sample the data every Wednesday – or the next trading day if Wednesday is a holiday.

We obtain data for seven indices: USA (SPX), Europe (SX5E), Germany (DAX), Switzerland (SMI), United Kingdom (UKX), Italy (MIB), Spain (IBEX). For each index, Table 4 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. (2016) for details on the weekly S&P 500 options. Finally, we only retain options whose prices are at least threefold the minimum tick size.

3.2 High-Frequency Equity-Index Futures Data

We obtain intra-day observations on the futures written on the underlying equity indices from TickData. We extract the futures price each minute, but our 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 a measure of the total quadratic variation of the log-price over the 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 jump variation (NJV), indicating the jump variation due to negative jumps. The measures are obtained following the procedure of Bollerslev and Todorov (2011) and Andersen et al. (2015b); see the Appendix for details. Table 5 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 only overlap with the Italian trading period by little more than two hours per day.⁴

Table 1 provides summary statistics for the daily measures, obtained from the high-frequency index returns. The vast divergence in the fortunes of the indices is striking. While the U.S. market experiences annual returns in excess of 4.5%, and the German index does slightly better, the Italian index drops by a full 10% annually over this eight year period, and the Spanish index is down an average of 4% per year. The indices experiencing negative excess returns generally have higher realized return variation measures – consistent with the so-called “leverage effect” – yet the average German volatility is also high and this index generates very attractive returns. The decomposition of the return variation stemming from large squared negative jumps versus the overall variation suggests that the U.S. index was less exposed to this type of negative shocks, while the differences across the European indices are minor. However, the relative jump counts are skewed by the fact that the U.S. markets are closed during the early trading in Europe, when many jumps may have materialized, but simply could not be observed in the U.S. The split into positive and negative return jumps reveal that the indices associated with the euro-zone had substantially more downward than upward jumps, possibly reflecting the impact of the sovereign debt crises. In contrast, the jump direction is nearly symmetric for the remaining indices. The statistic signaling slow updating of the index values (stale prices) indicates that the Swiss, and maybe the euro-zone ESTOXX, index may be affected by illiquidity, inducing a potential downward bias in the realized

⁴The 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.

variation measures, while the remaining indices are highly liquid. Nonetheless, the closeness of the return standard deviation measure obtained from the high-frequency and daily returns suggest that the impact of illiquidity on the measures is minimal.

	SP	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
\sqrt{RV}	22.49	26.94	24.63	19.73	23.20	26.36	25.80
\sqrt{TV}	21.79	25.59	23.27	18.46	22.22	25.20	24.40
JV/RV	0.06	0.10	0.11	0.12	0.08	0.09	0.11
NJV/JV	0.49	0.58	0.56	0.56	0.52	0.56	0.52
Average Daily Return (%)	4.58	-3.57	4.81	0.19	0.47	-9.99	-4.07
Std Daily Return (%)	22.49	27.60	25.83	20.05	23.77	27.76	27.08
% of Stale prices (%)	0.64	4.11	0.75	9.13	0.69	1.12	1.75

Table 1: **Summary Statistics for Equity-Index Futures.** The numbers are annualized and given in percentage form, except for the ratios in rows 3-4 and the last row dealing with stale prices. The realized variation measures (RV and TV) are computed from log returns within the trading day using the procedure detailed in the Appendix, then averaged and scaled to represent one calendar year. We report the square-root of these measures so they represent annualized standard deviation units. The daily standard deviation (std) is computed from close-to-close index return. The stale price statistics indicate the percentage of intervals for which neither the bid nor ask price changed across the five observations obtained by retaining the last quotes for each calendar minute in a given five-minute period.

Hence, the realized excess returns deviate greatly across the indices in our sample. At the same time, the indices respond roughly similarly to the major shocks during the financial and sovereign debt crises, as evidenced by the strong positive correlation between the daily equity-index returns displayed in Table 6 of the Appendix.⁵ In comparison, the observed discrepancies of the realized volatility measures across indices are less dramatic, albeit still highly statistically significant.

Our challenge is to provide a simple common framework for modeling the distinct return risks (future volatilities and jump realizations) and risk pricing (as reflected in the option prices) across this set of indices. These quantities jointly impact the individual equity and variance risk premiums over the sample whose realizations, as observed above, are dramatically different. In particular, a priori, it seems difficult to account for large realized risk premium differentials if risk pricing is

⁵The daily return correlations must be interpreted with caution, as they are not fully synchronized, especially for the U.S. and European indices. Effectively, the table provides lower bounds for the true (synchronized) correlations.

linked closely to volatility factors, which do not deviate sharply across the indices. Our modeling framework, developed below, introduces a new jump factor into the risk pricing for the equity-index options. This facilitates a more direct separation of volatility from risk pricing and enables the shape of the option surface to speak more cleanly to the (conditional) pricing of equity risk.

4 Model

We denote the generic equity market index by X . 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) . The corresponding (instantaneous) jump intensity, under the risk-neutral probability, also labeled the jump compensator, is $dt \otimes \nu_t^{\mathbb{Q}}(dx)$. The difference, $\tilde{\mu}^{\mathbb{Q}}(dt, dx) = \mu(dt, dx) - dt \nu_t^{\mathbb{Q}}(dx)$, constitutes the associated martingale jump measure.

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 is given by,

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

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 jump intensities are governed by the $c^-(t)$ and $c^+(t)$ coefficients which evolve as affine functions of the state vector,

$$\mathbf{c}^-(\mathbf{t}) = c_0^- + c_v^- V_{t-} + c_u^- U_{t-}, \quad \mathbf{c}^+(\mathbf{t}) = \mathbf{c}_0^+ + c_v^+ V_{t-} + c_u^+ U_{t-}.\tag{3}$$

This representation involves a large set of parameters. At the estimation stage, we zero out insignificant coefficients. Specifically, for the negative jump intensity, we fix c_0^- and c_v^- to zero and, for identification, normalize c_u^- to unity, so U equals the negative jump intensity, i.e., $c^-(t) = U_{t-}$. Also, we set $c_v^+ = c_u^+ = 0$, implying that the positive jump intensity is constant, $c^+(t) = c_0^+$. The terms relevant for our empirical implementation are printed in bold in equation (3).

To summarize, our jump modeling involves a number of novel features. First, the price jumps are exponentially distributed, unlike most prior studies which 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 squared price jumps impacting the factor dynamics in a manner reminiscent of discrete GARCH models. Third, the jump intensity is decoupled from the volatility. This is unlike most existing option pricing models in the literature, with the notable exception of Christoffersen et al. (2012) and Li and Zinna (2015). Nonetheless, as stressed by Andersen et al. (2015b), model (1) still belongs to the affine class of models of Duffie et al. (2000). For future reference, we label our two-factor affine model, including the negative jump intensity U_t , the *2FU* model.

5 Estimation Procedure

We follow the estimation and inference procedures developed in Andersen et al. (2015a). The option prices are converted into the 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 option price is given by $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$. Estimation of the parameter vector and the period-by-period realization of the state vector now proceeds by minimizing the distance between the observed and model-implied BSIV across the full sample in a metric that also penalizes for the discrepancies between the inferred spot volatilities and those estimated (in a model-free way) from high-frequency return observations on the underlying asset, $\sqrt{\widehat{V}_t^n}$. The latter are obtained from five-minute returns over a three-hour window prior to the close of the trading day using the so-called truncated realized volatility estimator, as implemented in Andersen et al. (2015b), see the Appendix for further details. 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 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 $\mathbf{S}_t = (V_t, U_t)$ from the following optimization problem,

$$\left(\{\widehat{\mathbf{S}}_t\}_{t=1}^T, \widehat{\theta} \right) = \underset{\{\mathbf{S}_t\}_{t=1}^T, \theta}{\operatorname{argmin}} \sum_{t=1}^T \left\{ \sum_{\tau_j, k_j} \frac{(\kappa_{t, k_j, \tau_j} - \kappa_{k_j, \tau_j}(\mathbf{S}_t, \theta))^2}{N_t} + \frac{\xi_n}{N_t} \frac{(\sqrt{\widehat{V}_t^n} - \sqrt{V_t})^2}{\widehat{V}_t^n/2} \right\}, \quad (4)$$

where the penalty for the deviation between the realized and model-implied spot volatility is given by $\xi_n > 0$ and the superscript n denotes the number of high-frequency returns exploited by the spot return variance estimator, \widehat{V}_t^n . For our implementation with a given fixed n , we set $\xi_n = 0.05$, as in Andersen et al. (2015b). Moreover, 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 sampling widely across the full sample period. The shape of the surface varies dramatically across 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 (4). Thus, we have daily estimates for the state realizations available, even if full-fledged estimation is performed only for weekly data.

6 Option Factors

Given our limited sample period and the lower number of observations available for some of the European indices relative to the S&P 500 illustration in Section 2, our purpose is not to obtain a perfect option pricing model, but rather to settle on a specification that captures the salient features across all indices in a robust manner. Thus, our specification involves only a single volatility factor and the dynamics of the negative jump intensity factor is pared down relative to the model in 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. We start by illustrating the extent to which the jump intensity series for the S&P 500 index extracted using the current model is similar to the one obtained from the more elaborate specification in Andersen et al. (2015b).

6.1 Comparison with a More Elaborate Asset Pricing Specification

As discussed earlier, our model is parsimonious, and in particular it contains only one volatility factor. One potential concern is that the relatively constrained modeling of the volatility process may induce excessive bias into 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 volatility series, clearly outperforming more standard and equally heavily parameterized representations.

The sample exploited in this paper is shorter than in Andersen et al. (2015b), due to the synchronization with the data available for the European markets. However, the current sample 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 July 23, 2013. Hence, the two series of option-implied jump factors overlap over the period January 2007 till July 2013.

Figure 5 depicts the jump factor extracted from the two separate models, estimated from partially overlapping periods and option samples. In spite of these differences, the models deliver remarkably similar time series paths for the jump factor, corroborating the assertion that the dynamic specification of either model enables robust identification of this factor.




Figure 5: **Left Jump Tail Factor Comparison.** The figure depicts the daily option-implied left jump tail factor from the models (1) and the one in Andersen et al. (2015b). 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. The correlation between the two extracted jump factors in the period of overlap is 0.987.

6.2 Country-by-Country Factor Realizations

We now provide a more detailed account of the implied factors extracted from the various equity-index option markets based on our $2FU$ (two factors, including U) model (1). We first compare the implied spot variance for each of the individual equity indices enumerated in Section 3 to the spot variance extracted from the S&P 500 options, depicted in Section 2. Figure 6 plots the extracted spot 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 some 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 only deviates for a couple of episodes during the fixed Swiss franc-euro exchange rate policy implemented in September 2011 and lasting through the remainder of our sample. In the former case, the correlation between the volatilities is about 98% and in the latter around 95%. In contrast, the 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 significantly more than the S&P 500, and a positive gap remains for the remainder of the sample, albeit to a varying degree. For the broader euro-zone ESTOXX index, the same effect is clearly visible and originates with the first phase of the European debt crisis. Thus, while the volatility patterns were strikingly similar for some of the primary equity indices in North America and Europe through the financial crisis, the impact of the sovereign debt crisis is clearly heterogeneous. Moreover, the relative effect across the national indices appears consistent with the perceived sensitivity of the respective economies to the European crisis. This is especially transparent 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 the latter indices attain a plateau well above the others ever since the first signs of 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.

Next, Figure 7 depicts the option-implied (risk-neutral) negative jump variation for these indices along with the corresponding quantity for the S&P 500. Formally, the risk-neutral negative jump variation is defined as $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx)$. In the $2FU$ model it equals $\frac{2}{\lambda^2} U_t$, so these jump variation series are directly proportional to the negative jump intensity factor, U . Consequently, the relative variation in the series reflects the corresponding variation in the extracted negative jump intensity state variables for the individual indices.

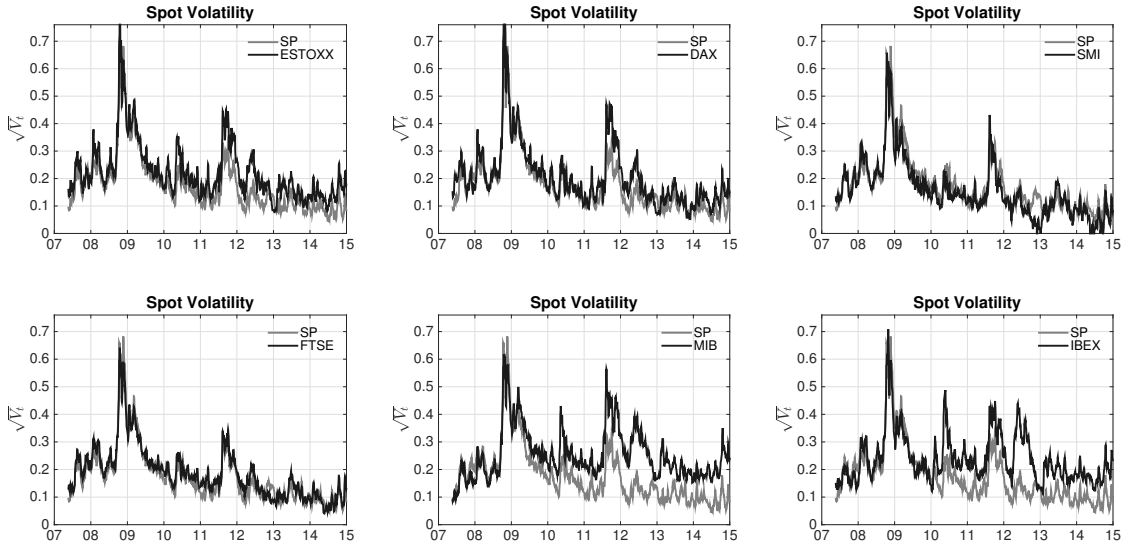


Figure 6: **Volatility Factor Comparison.** For each option-implied spot variance daily factor, we report the trailing five-day moving average of $\sqrt{V_t}$. The pairwise correlation between the volatility factor of the S&P 500 and each of the European indices is as follows: SP-ESTOXX: 0.948; SP-DAX: 0.939; SP-SMI: 0.953; SP-FTSE: 0.979; SP-MIB: 0.750; and SP-IBEX: 0.770.

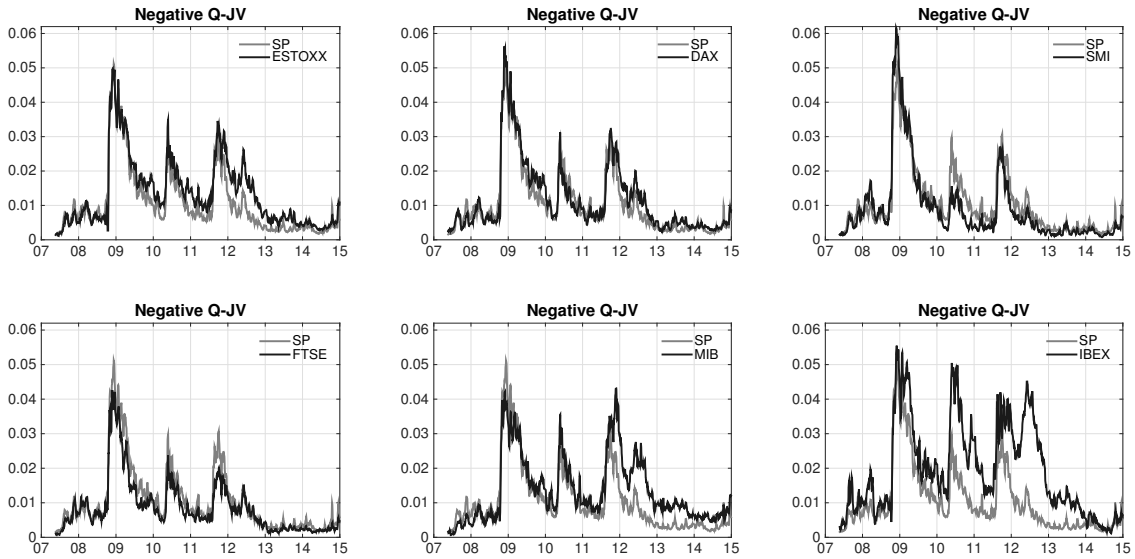


Figure 7: **U Factor Comparison.** For each option-implied negative jump intensity factor, we report the trailing five-day moving average of the implied risk-neutral jump variation $\int_{x<0} x^2 v_t^{\mathbb{Q}}(dx) \equiv \frac{2}{\lambda^2} U_t$. The pairwise correlation between the jump factor of the S&P 500 and each of the European indices is as follows: SP-ESTOXX: 0.949; SP-DAX: 0.970; SP-SMI: 0.951; SP-FTSE: 0.981; SP-MIB: 0.854; and SP-IBEX: 0.798.

Qualitatively, the pattern is similar to the one observed for the volatility factors. For each index, the volatility and jump factors are highly correlated. Nonetheless, as for the U.S., the relative size of the spikes in the jump intensity versus volatility varies substantially, with the European crisis inducing a stronger surge in the jump intensity relative to (diffusive) volatility. Again, the U.K. and U.S. series evolve in near unison and display a correlation of over 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. before and during the financial crisis and then below after the summer of 2009. We also observe strong coherence between the S&P 500 series and the ESTOXX and DAX series up through the financial crisis and then a relative elevation in the latter from the summer of 2009 and onwards, with the effect being notably more pronounced for ESTOXX than DAX, again suggesting a smaller economic exposure of Germany to the debt crisis than for the broader euro-zone. The most striking contrast occurs for the two Southern European indices, however. 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 two 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 adds further credence to the robustness of our methodology in extracting the salient pricing features from the option surfaces.

7 Option-Based Prediction of Future Risk and Risk Premium

We now explore the ability of the option-implied factors, spot variance and left jump intensity, to forecast the (realized) future return variation – defined as the sum of the squared high-frequency equity-index futures returns – and the (realized) equity excess return. The former signifies whether the factors are associated with market-wide risk as captured by equity volatility and jump activity, while the latter speaks to their predictive power for the equity risk premium. A third subsection compares the predictive performance for the equity risk premium to the one obtained from the variance risk premium, which has been stressed as an important explanatory variable in the recent

literature.

We rely on standard predictive regressions to assess the forecast power of the state variables. Similar to the case of the S&P 500 index, to avoid look ahead bias, we re-estimate the parameters of the risk-neutral $2FU$ model for all indices using only data from 2007. We then implement the predictive regressions conditional on these predetermined model parameters for the “out-of-sample” period 2008-2014.⁶ Specifically, we extract the state vector realizations from the objective function (4) by optimizing over S_t for each date t over the period 2008-2014, with θ fixed at the parameter vector estimated from the 2007 data.

7.1 Predicting Equity Risk Premium

This section focuses on the evidence for predictability of the equity risk premium. We regress the future excess returns for each of our equity indices on their option-implied state variables. For all indices, the two state variables are correlated so, for ease of interpretation, we supplement the spot variance factor, V , by the component of the left jump factor that is orthogonal to the spot variance, denoted \tilde{U} , as the second explanatory variable. This approach will ascribe all explanatory power that stem 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 regression that is orthogonal to the variance factor, i.e., the incremental information in the tail factor.

Our regression takes the following form for each index and $t = 1, \dots, T - h$,

$$r_{t,t+h} = \log(X_{t+h}) - \log(X_t) = c_{0,h} + c_{v,h} \cdot V_t + c_{u,h} \cdot \tilde{U}_t + \epsilon_{t,h}, \quad (5)$$

where h is the horizon (in days), $r_{t,t+h}$ denotes the h -day continuously-compounded return, and \tilde{U}_t is the option-implied left jump intensity orthogonalized with respect to the spot variance.

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. We compute robust Newey-West standard errors using a lag length that is twice the number of weeks within the forecast horizon. Finally, given the short sample period and the variability in liquidity for some of our index options, the results can be sensitive to outliers. The more extreme observations may be genuine, but can also arise from data errors, non-synchronous observations, or occasional poor identification of the factor realizations. Thus, for robustness, we winsorize all explanatory variables in the subsequent

⁶The one-year estimation period renders the point estimates less accurate than those obtained for the full sample. Nonetheless, the general shape of the volatility surface is quite stable, enabling the short sample to capture the salient features of the risk-neutral dynamics. We have confirmed that the extracted realizations of the state vector for 2008-2014 are qualitatively similar whether based on parameter estimates for 2007 or the full sample. The estimates for the 2007 sample are available in a Supplementary Appendix, available upon request.

predictive regressions at the 98 percent level, thus limiting the influence of the 1% extreme negative and positive observations. Importantly, we do not exclude the corresponding daily return or return variation from the cumulative measures appearing as left-hand-side variables, i.e., the extreme return or volatility realizations are included in the definition of the multi-horizon excess returns and return variation measures. Finally, the inference is based on Newey-West robust standard errors, taking into account the estimation errors in any potential first-stage regression used to normalize the regressors.

The regression results are, qualitatively, similar across the indices. For example, the top left panel of Figure 8 plots the t-statistics for the ESTOXX regression slopes in equation (5), while the top right panel displays the corresponding regression R^2 . Consistent with the hypothesis of a slowly moving equity risk premium, the predictive power of the regression rises steadily over time. We note that the unpredictable return component seems to dominate at the weekly horizons, but as the random noise innovations cancel over time, the predictable component emerges clearly for the longer horizons. Within the four month mark, the R^2 surpasses 10% and exceeds 15% after six months. It is further evident that the explanatory power stems primarily from the jump intensity factor, while the variance factor is insignificant for all horizons beyond one week. 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. Of course, this is also consistent with our findings for the S&P 500 index in Section 2 as well as Andersen et al. (2015b).

Turning to the remaining indices in Figure 8, we observe strikingly similar features across them all. Apart from Swiss SMI index, we find the residual left tail factor to attain significance at the (one-sided) 2.5% level around the three month mark. The degree of explained variation is consistently between 15-20%, again excluding the Swiss index. And, uniformly, the variance factor is bereft of explanatory power in these regressions. These findings are quite remarkable given the extremely different equity returns the indices have offered over the sample, and the highly diverse exposures they display vis-a-vis the European debt crisis. Hence, the finding that the spot variance has no systematic relation to the future excess returns is robust across this sample period. Moreover, this conclusion is in line with an extensive time series oriented 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.

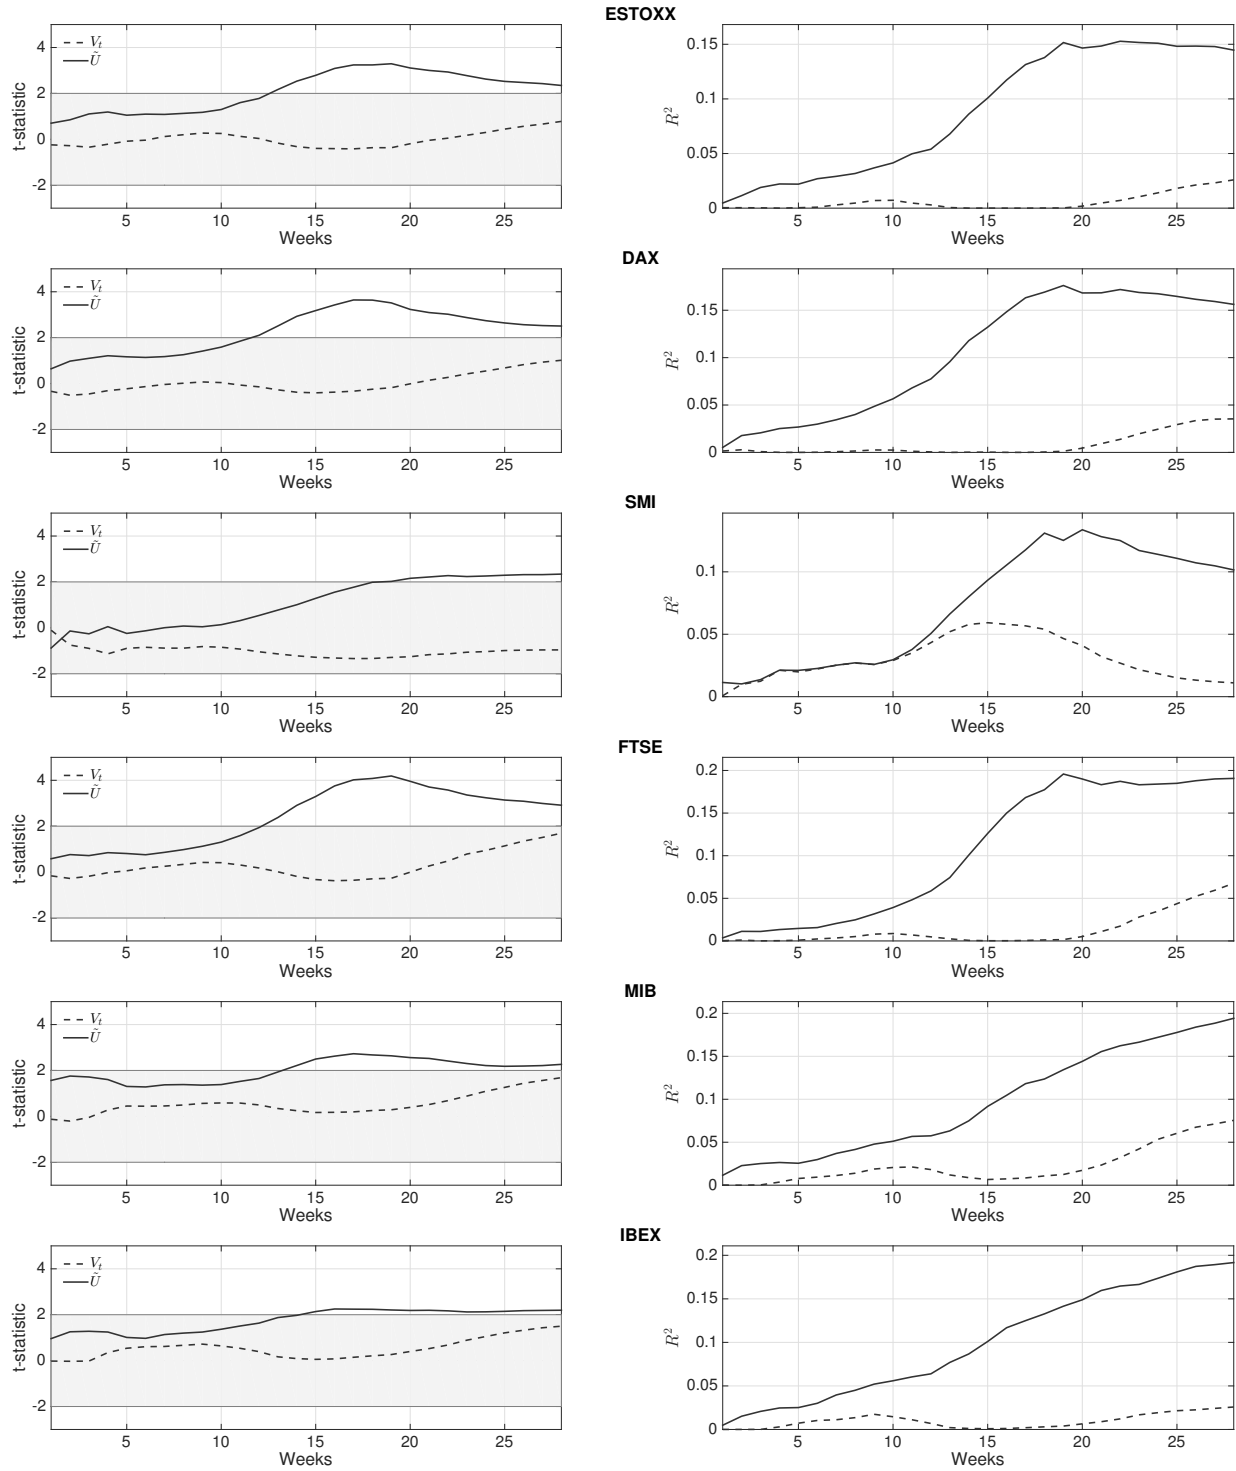


Figure 8: **Predictive Regressions for Equity Excess Returns.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

7.2 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 again consider regressions of the form (5), but the dependent variable is now a measure of the future realized return variation over the time interval $[t, t + h]$, $RV_{t,t+h}$. The latter is constructed from the 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, so they provide good proxies for the risk associated with exposure to the market index, see, e.g., Andersen et al. (2003). The regressions take the form,

$$RV_{t,t+h} = k_{0,h} + k_{v,h} \cdot V_t + k_{u,h} \cdot \tilde{U}_t + \epsilon_{t,h}, \quad (6)$$

The left panels of Figure 9 reveal, in line with the prior S&P 500 evidence, that the left jump intensity factor has no explanatory power for the ex-post realized return variation for any of the indices. 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 complete absence of predictive power in the jump factor is striking. This is the component that encompasses the relevant information for predicting the equity returns, but it is entirely unrelated to return volatility.

To summarize, we document a clear empirical separation between the determinants of the equity risk premium and market risks: the latter are well captured by the level of market volatility while the former is driven by the component of the option-implied risk-neutral left jump intensity not spanned linearly by market volatility. The finding is consistent across all indices in our analysis.

7.3 Comparison with the Variance Risk Premium

Our predictive results can rationalize the US and international evidence of Bollerslev et al. (2009) and Bollerslev et al. (2014), respectively, that the country-specific variance risk premium has predictive power for that country's future excess returns.⁷ The variance risk premium is the compensation demanded by investors for bearing variance risk. Formally, it is defined as the gap between the conditional risk-neutral and statistical expectation of future return variation. The former can be easily measured in a model-free way using a portfolio of close-to-maturity options and it corresponds to

⁷Our findings are also related to the broader extensive literature documenting return predictability for international equity indices, see e.g., Harvey (1991), Bekaert and Hodrick (1992), Campbell and Hamao (1992), Ferson and Harvey (1993) and Hjalmarsson (2010).

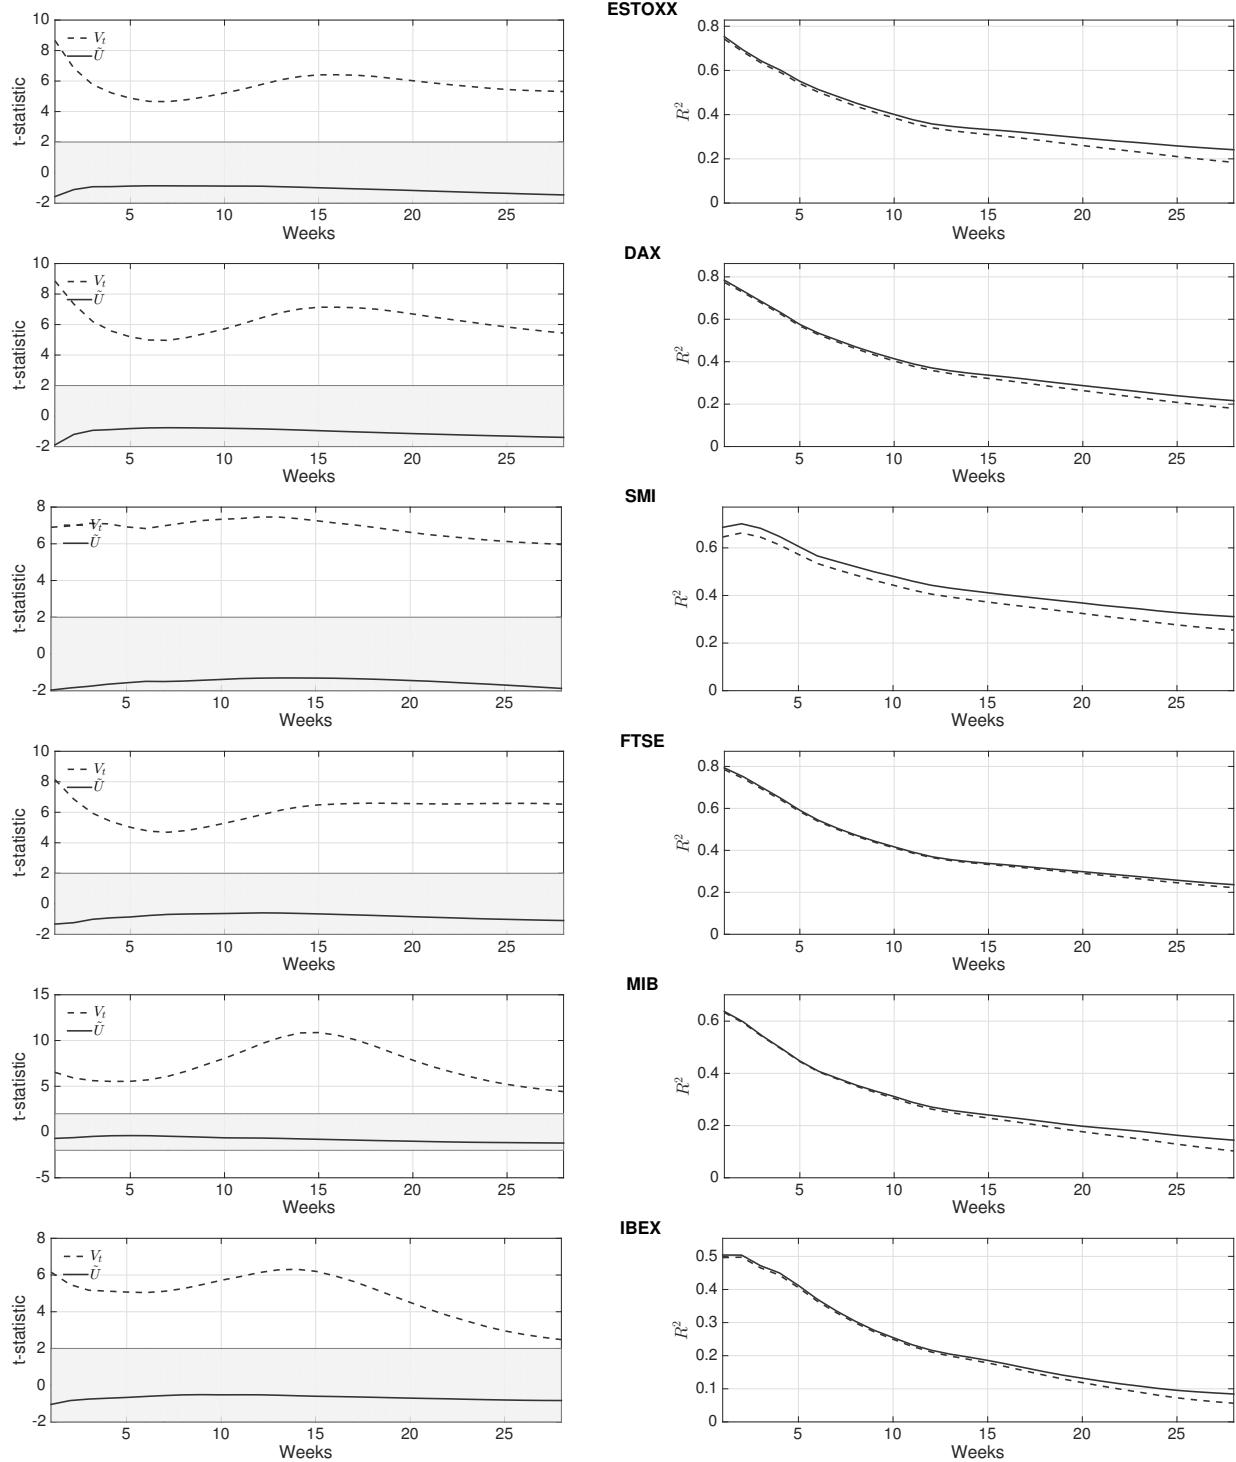


Figure 9: **Predictive Regressions for Return Variation.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

the value of the VIX index. Since tail risk is part of the total return variation, our U factor forms part of the VIX index. In fact, within our risk-neutral model, the VIX index is a simple affine function of the two option factors, U and V . Therefore, a natural question that arises in light of our findings above is whether the predictive ability of the variance risk premium is solely due to the “pure-tail” factor we discovered here, i.e., \tilde{U} .

To address this question, we will explore predictive regressions for the equity premium including both the tail factor and the variance risk premium (VRP) as explanatory variables. As in Section 2, we estimate the expected future return variation using the model of Corsi (2009). However, in contrast to the U.S. case, where we rely on the VIX measure provided by the CBOE, we do not use official exchange measures for the risk-neutral return variation. Instead, this component of the variance risk premium is computed following the procedure of Carr and Wu (2009).⁸

Given the strong correlation among the two variables, we focus on the auxiliary explanatory power provided by the variance risk premium in the following specification,

$$r_{t,t+h} = \log(X_{t+h}) - \log(X_t) = c_{0,h} + c_{u,h} \cdot \tilde{U}_t + c_{p,h} \cdot \widetilde{VRP}_t + \epsilon_{t,h}, \quad (7)$$

where \widetilde{VRP}_t denotes the (estimated) variance risk premium at the end of trading day t , orthogonalized with respect to the tail factor developed in the preceding sections.

The results are summarized in Figure 10. The qualitative features are remarkably similar to those of Figure 8. In fact, the overall regression R^2 is now, if anything, lower than before and the significance of the jump factor remains strong for the longer return horizons. Finally, there is no systematic evidence of auxiliary predictive power of the variance risk premium over-and-above what is captured by the jump factor. In summary, we conclude that the jump factor is the driving force behind the predictive power of the variance risk premium for future equity excess returns. We strongly reject the reverse hypothesis that the variance risk premium is a superior predictor relative to our negative jump intensity factor.

8 Downside Jump Risk and Risk Premium Linkages

We have identified the negative jump factor as a primary indicator of the future equity risk premium. Since this factor represents the risk-neutral negative jump intensity, it may increase through two

⁸Specifically, 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 constant and equal to the implied volatility at the lowest (highest) available strike level. Finally, we compute the VIX index for both maturities and linearly interpolate to obtain the VIX index corresponding to 30 calendar days.

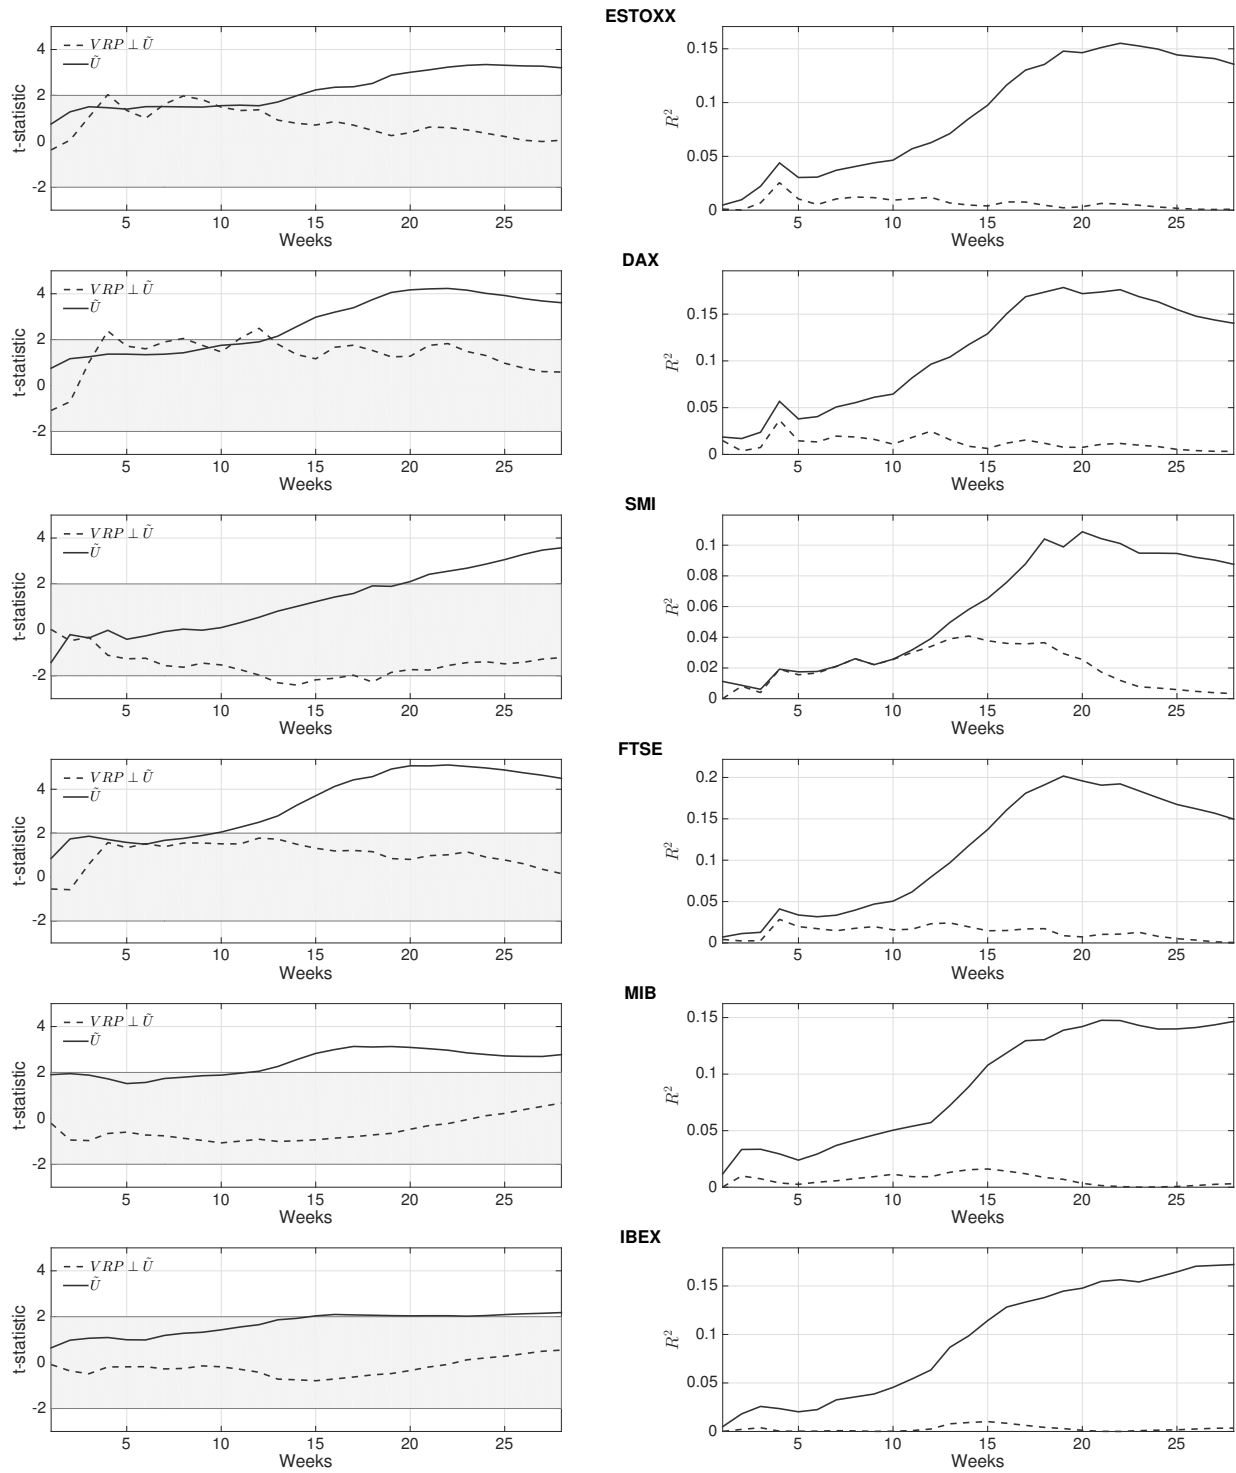


Figure 10: **Predictive Regressions of Excess Returns.** Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the auxiliary explanatory power afforded by the VRP.

separate channels. It may be due to a rise in the expected left jump intensity under the objective measure or, alternatively, it may reflect an elevation of the negative jump risk premium, i.e., the gap between the negative jump intensity under the risk-neutral (\mathbb{Q}) and the objective (\mathbb{P}) measure.

Given the critical role of the jump factor, it is of interest to gauge the extent to which each of the channels identified above contribute to fluctuations in the option-implied negative jump intensity and to the co-movements of these measures across the international markets.

While risk-neutral expectations of future jump risk are easy to generate (and are readily provided by our parameter and factor estimates), it is rather difficult to compute reliable estimates for their counterparts under the statistical probability, especially given our short and unusually turbulent sample. Large jumps are influential in evaluating the jump variation, but they are relatively rare and cluster, leading to imprecise small-sample inference. Under such circumstances, one may turn to a tight parametric model that facilitates identification of the link between the jump variation and the underlying state variables. However, we have, so far, avoided imposing parametric structures on the \mathbb{P} distribution, as we instead rely on inference under the \mathbb{Q} measure, which embeds the market's assessment of the state variables and their (risk-neutral) dynamics into the pricing of the option surface. There is no corresponding source of rich information regarding the price dynamics under the objective measure, so those parametric assumptions on that dimension are much harder to assess.

In view of these difficulties, we instead explore the observed realized negative jump variation across different indices and use the findings to learn about the evolution of their relative jump risk premiums, $NJV_{t,t+h}$. These realized jump variation measures are given by,

$$NJV_{t,t+h} = \frac{1}{h} \left(\sum_{s \in [t,t+h]} (\Delta X_s)^2 1_{\{\Delta X_s < 0\}} \right). \quad (8)$$

Although large idiosyncratic innovations renders estimation of the mean and standard deviation of the jump variation difficult, we may infer a great deal about the relative exposure of the various equity indices to shocks that induce actual jumps in asset returns. This may be achieved, throughout this eventful period for the global markets, by measuring whether the realized jump activity for the indices display similar or divergent jump behavior. To the extent their responses are well aligned, it is natural to infer that the (a priori) expected jump variation measures also were highly correlated.⁹

Nonetheless, there are additional complications to confront before undertaking this type of analysis. One stems from a review of our prior findings regarding the option-implied (priced)

⁹This follows if agents rationally update their expectations about the relative jump exposures of the indices to reflect the observations over time.

negative jump variation. In Section 6.2, we document a strong qualitative correspondence between the negative jump factor for the U.S., U.K., and Germany over the period 2007-2009. However, thereafter, during 2010-2014, the German jump factor rises above the others in periods surrounding the main phases of the sovereign debt crisis. Similarly, the Italian jump factor is aligned with these three factors through the great recession, but elevates more than the others from the onset of the sovereign debt problems. Even the Spanish factor is roughly proportional with the above factors through 2009, although it clearly is more volatile. In contrast, after the debt crisis emerges in 2010, the Spanish jump factor fluctuates violently throughout the remainder of the sample and frequently attains truly extreme levels compared to the other indices. This suggests the presence of a structural break in the relationship between the equity indices in 2010, with the jump exposures diverging significantly from the patterns observed over the first three years of the sample.

A second complication is the different currency denominations of the indices. While large currency movements may not have been anticipated a priori, we expect economic shocks, that induce differentially sized jumps in the equity indices, to induce concurrent jumps in the exchange rates. To avoid such confounding effects, we exploit the fact that three of them are denominated in the identical (euro) currency, namely the German (DAX), the Italian (MIB) and the Spanish (IBEX) indices. As just noted, these indices also display a highly divergent dynamic in the pricing of jumps across the sample, rendering them excellent candidates for an exploration of potential shifts in the relative jump risk premiums over time.

8.1 Realized Jump Risk Co-Movements in the Euro-Zone

In light of the complications discussed above, this section focuses on the three euro-denominated indices, and we conduct a separate analysis for the first and second part of the sample to allow for a potential break in the objective jump exposure of the indices.

To assess their relative exposure to the jump factor, we first extract a series of index return jumps from each individual market, using five-minute index futures data and the jump identification techniques detailed in the Appendix. We restrict the sample to the common trading hours across the exchanges, so that we can monitor their simultaneous response to events that generate return jumps. Specifically, we classify the jumps for a given index as common or country-specific depending on whether the jump occurrence is observed simultaneously with jumps in *both* of the other indices or not. The arrival of common jumps enables us to assess the relative sensitivity of the indices to global events that have a significant affect on all of them, while the country-specific jumps are indicative of residual or more idiosyncratic jump activity.

Figure 11 displays the 344 common jumps for the three indices plotted pairwise against each

other over the period 2007-2009, encompassing the great recession but not the sovereign debt crisis. The plots reveal a remarkable similarity in the jump response across the three markets. The regression lines all have a slope close to, and insignificantly different from unity, suggesting that they move one-for-one in response to the news generated around the financial crisis. The associated R^2 values range from 0.88 to 0.94. In brief, the common jumps appear to reflect a similar exposure of all three indices to this type of event.




Figure 11: **Scatter of Common Market Jumps in Germany, Italy and Spain.** The straight lines on the plots represent the fit from a linear regression (without an intercept) of the jumps on the Y-axis on the jumps on the X-axis.

The country-specific jumps cannot be analyzed in the same manner as they, by definition, do not occur simultaneously. Instead, we assess the residual jump activity by computing the country-specific monthly realized negative and positive jump variation across the sample period. These realized measures are depicted in Figure 12. Focusing on the 2007-2009 period, we do not observe any major differences in the series. The most noteworthy discrepancy is the limited degree of positive jump variation in the DAX index. However, on the downside, the DAX jump variation spikes along with the IBEX at the beginning of 2008 and during the depth of the financial crises, while the MIB response is much smaller in the first instance but more dramatic in the second. Finally, after October 2008, the DAX appears less susceptible to large negative idiosyncratic shocks than the other two indices. Nonetheless, overall, for the period from 2007 through the main part of the financial crisis, there is no evidence of systematic discrepancies in the realized negative jump exposures of these equity indices. These conclusions are corroborated by the summary statistics for the monthly country-specific jump variation measures in Table 2. In particular, it is evident that the correlation of the DAX measures with the other two are much stronger for the negative than the positive jump variation.¹⁰

¹⁰This is consistent with evidence from lower frequency data, see e.g., Longin and Solnik (2001), Ang and Chen

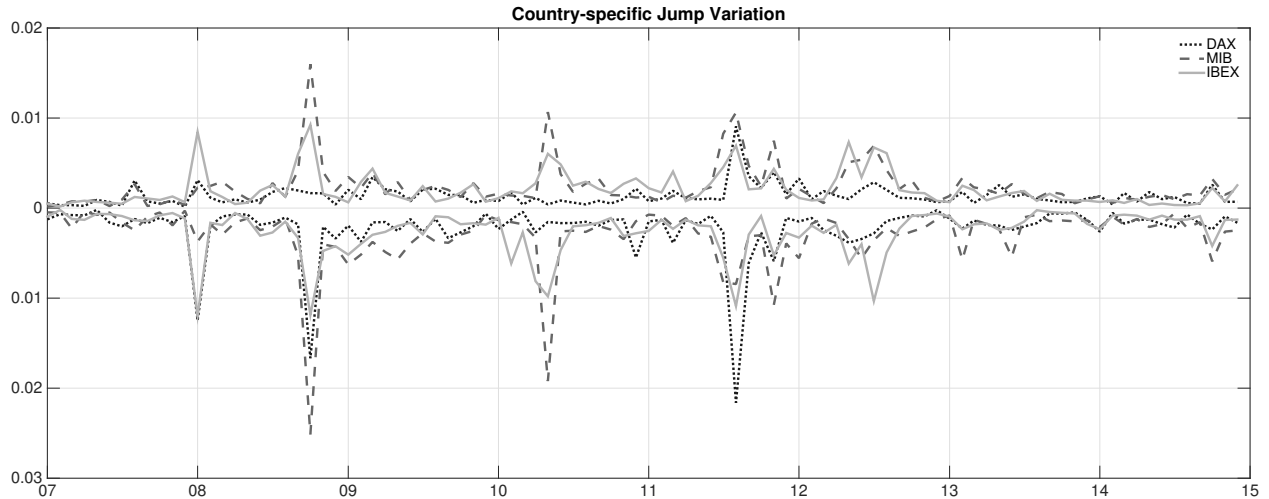


Figure 12: **Country-specific Jump Variation in Germany, Italy and Spain.** The jump variation due to the positive and negative jumps is above and below, respectively, the zero line. The jump variations are computed over periods of months from the detected jumps over these intervals. The jump variations are reported in annualized decimal units.

Figure 13 provides the corresponding scatter plots for the 547 common jumps over the period 2010-2014. The jump realizations continue to be approximately linearly aligned. However, the relative amplitudes have shifted significantly. Now, the DAX index responses are more muted than for the Southern European counterparts. The slopes for the middle and right panels in the figure are 0.80 and 0.75, respectively, while the slope in the left panel remains indistinguishable from unity. Hence, while the Italian and Spanish markets continue to display similar responses to common shocks, the German exposure has declined in relative terms. This is consistent with the sovereign debt crisis having a disproportional impact on Italy and Spain compared to Germany. Not surprisingly, the R^2 values for the DAX plots, at 0.87 and 0.84, have dropped relative to before, while the value for the IBEX-MIB plot is 0.94 and, if anything, has increased relative to the financial crisis, speaking to a great deal of commonality in the jump exposures for Italy and Spain during the debt crisis.

Inspection of Figure 12 suggests a fair degree of commonality in the country-specific jump variation measures over the second part of the sample. However, the German country-specific jump exposure now appears decidedly lower than for the other two indices. In fact, apart from the fairly dramatic spike associated with the onset of the second phase of the sovereign debt crisis, the

(2002), Poon et al. (2004), Bae et al. (2003) and more recent evidence from high-frequency stock data, see, e.g., Bollerslev et al. (2013).




Figure 13: **Scatter of Common Market Jumps in Germany, Italy and Spain.** The straight lines on the plots represent the fit from a linear regression (without an intercept) of the jumps on the Y-axis on the jumps on the X-axis.

realized German jump variation is almost uniformly lower than the others. In comparison, Spain and Italy alternate in terms of displaying the largest realized jump exposure, with no evidence of one being, on average, more susceptible to large idiosyncratic return shocks than the other. These conclusions are confirmed by Table 3. Over this part of the sample, the German negative jump variation is less correlated with those in Italy and Spain than before, and there is no longer an asymmetry between the correlation of the positive and negative jump measures for the DAX. On the other hand, the correlation between the MIB and IBEX remains high for both the negative and positive jump measures. In short, the DAX displays a substantially lower degree of exposure to both large common shocks and more idiosyncratic events in the latter part of our sample, while there is little to separate the objective jump exposures among MIB and IBEX.¹¹

8.2 Implications for International Jump Risk Premium

Overall, our nonparametric analysis of the jump risks across Germany, Italy and Spain reveal strong linkages, but also a distinct shift in the relative exposures following the financial crisis. By combining these findings for the *realized* negative jump variation with those for the *priced* negative jump variation in Section 6.2, as depicted in Figure 7, we can make an informed assessment of the relative negative jump risk premium across the indices. Over the first part of our sample, the risk-neutral negative jump variation of the Spanish index is, on average, more elevated than the Italian, German and U.S. indices and it clearly displays more abrupt fluctuations. This contrasts sharply

¹¹The number of country-specific jumps are spread fairly evenly across the sample period, albeit slightly tilted towards the downside. Compared to the total of 891 common jumps, the number of positive and negative country-specific jumps in Germany, Italy and Spain are 1319, 1166 and 1213, versus 1546, 1364 and 1503, respectively.

	Mean ($\times 10^2$)	StD ($\times 10^2$)	Correlation		
Positive Jump Variation					
DAX	0.13	0.09	1	0.35	0.51
MIB	0.21	0.26		1	0.72
IBEX	0.18	0.21			1
Negative Jump Variation					
DAX	0.23	0.32	1	0.80	0.91
MIB	0.32	0.41		1	0.73
IBEX	0.25	0.26			1

Table 2: **Country-specific Jump Variation; 2007-2009.** Summary statistics. Mean and standard deviation are reported in annualized decimal units.

	Mean ($\times 10^2$)	StD ($\times 10^2$)	Correlation		
Positive Jump Variation					
DAX	0.14	0.13	1	0.54	0.42
MIB	0.25	0.22		1	0.79
IBEX	0.22	0.17			1
Negative Jump Variation					
DAX	0.22	0.28	1	0.33	0.54
MIB	0.29	0.29		1	0.65
IBEX	0.26	0.24			1

Table 3: **Country-specific Jump Variation; 2010-2014.** Summary statistics. Mean and standard deviation are reported in annualized decimal units.

with the actual realized jump exposure that is near indistinguishable across the three European indices over this period. We infer that the Spanish negative jump risk premium was larger and more volatile than the corresponding premium in Germany and Italy. One may conjecture that the

rising tensions in housing and mortgage markets, that preceded the financial crisis, were taking a disproportionate toll on the risk attitudes towards the Spanish economy which was fueled by a large real estate boom. In contrast, if anything, the Italian risk premium for negative jump risk appears smaller than in Germany during the height of the financial crisis as the spike in the Italian risk-neutral negative jump variation in Figure 7 is less elevated. This may reflect a perception of a relatively lower vulnerability of the Italian economy to the sharp contraction in the financial sector.




Figure 14: **Option-Implied Factor Spreads in the Euro-Zone.**

After 2010, the picture shifts dramatically. From Figure 7, it is evident that the Italian priced negative jump variation grows increasingly volatile and generally exceeds the one in Germany, especially in the wake of the second phase of the sovereign debt crisis in late 2011. As noted previously, however, the impact on the corresponding Spanish series is even more drastic. Given the near identical realized negative jump variation in Spain and Italy over this period, documented in Figures 12–13 and Table 3, we conclude that the Spanish premium for negative jump risk exceeds the Italian by a substantial margin. One likely explanation is the threat of Spain dispensing with the euro currency and joining a set of other nations in adopting a Southern euro in its place. Such scenarios may have been deemed less plausible in the case of Italy. For Germany, the priced negative jump risk remained fairly close to the U.S. level, indicating a much more modest elevation in the negative jump risk premium for the DAX during the European debt crisis.

9 Conclusion

This paper applies the option pricing approach of Andersen et al. (2015a) to a number of international equity indices, including the US 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, variance and jump tail risk premiums, and a spot variance factor which is a potent predictor of the

actual future return variation. Standard approaches exploiting only volatility factors miss the equity risk premium information in the option surface insofar as the volatility factors do not span the “pure tail factor,” which is the one embedding the predictive content for the equity risk premium.

For all major indices, the evolution and pricing of the financial market risks appear consistent. Nonetheless, important deviations appear around the sovereign debt crises in Europe, when the option surface associated with the indices denominated in the euro currency show a varying degree of elevated volatility and downside tail risk. Germany displays only a temporary degree of risk elevation, while the Italian and Spanish indices are dramatically affected. And even for these two countries, we observe significant variation, as the elevation in the Italian left tail factor is more muted and reflects mostly the increasingly volatile financial market conditions experienced following the onset of the sovereign debt crisis. In contrast, the Spanish tail factor is exceedingly elevated and only returns to more normal levels years after the crises. As such, the Spanish risk pricing is clearly unique within the set of indices explored.

Given our specific findings regarding the relative risk pricing in Italy and Spain, it will be of interest, in future work, to associate the inferred downside tail risk premium in these countries with the actual sequence of economic events affecting them during this period. Corresponding studies relating the relative risk pricing across equity indices denominated in different currencies will require explicit consideration of currency risk in the analysis. 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., 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 (2016). Pricing Short-Term Market Risk: Evidence from Weekly Options. *Journal of Finance*, forthcoming.
- Ang, A. and J. Chen (2002). Asymmetric Correlations of Equity Portfolios. *Journal of Financial Economics* 63, 443–494.
- Bae, K., G. Karolyi, and R. Stulz (2003). A New Approach to Measuring Financial Contagion. *Review of Financial Studies* 16, 717–763.
- 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 S. Li (2013). Jump Tails, Extreme Dependencies, and the Distribution of Stock Returns. *Journal of Econometrics* 172, 307–324.
- 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 Finance* 47, 43–69.
- Carr, P. and L. Wu (2009). Variance Risk Premia. *Review of Financial Studies* 22, 1311–1341.
- Christoffersen, P., K. Jacobs, and C. Ornthalalai (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.
- 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.
- Kou, S. (2002). A Jump- Diffusion Model for Option Pricing. *Management Science* 48, 1086–1101.
- Li, J. and G. Zinna (2015). The Variance Risk Premium: Components, Term Structures, and Stock Return Predictability. Working paper.
- Longin, F. and B. Solnik (2001). Extreme Correlation of International Equity Markets. *Journal of Finance* 56, 649–676.
- Merton, R. (1976). Option Pricing when Underlying Asset Returns are Discontinuous. *Journal of Financial Economics* 3, 125–144.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Poon, S., M. Rockinger, and J. Tawn (2004). Extreme Value Dependence in Financial Markets: Diagnostics, Models, and Financial Implications. *Review of Financial Studies* 17, 581–610.
- 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.

10 Appendix

10.1 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 4: **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).

Each exchange provides detailed information about the trading specifications of each option contract:

- SPX: http://www.cboe.com/products/indexopts/spx_spec.aspx,
- EURO STOXX 500: <http://www.eurexchange.com/exchange-en/products/idx/stx/blc/EURO-STOXX-50--Index-Options/19066>,
- DAX: <http://www.eurexchange.com/exchange-en/products/idx/dax/DAX--Options/17252>,
- SMI: <http://www.eurexchange.com/exchange-en/products/idx/smi/SMI--Options/19508>,
- FTSE 100: <https://www.theice.com/products/38716770/FTSE-100-Index-Option>,
- FTSE MIB: <http://www.borsaitaliana.it/derivati/specifichecontrattuali/ftsemiboptions.en.htm>,
- IBEX 35: <http://www.meff.es/aspx/Comun/Pagina.aspx?l1=Financiero&f=OpcionesIBEX35&id=ing>.

10.2 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 5: **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).

10.3 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:

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

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

2. Bipower Variation:

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

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 $|\Delta_i^n Z| >$.

4. Truncated Variation and Jump Variation:

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

5. 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 (4), we take into account the following:

- $\widehat{V}_t^n = \frac{1}{h} TV_{t,t+h}$, where h 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.

10.4 Equity-Index Return Correlations

	SP	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
SP	1.00	0.80	0.79	0.70	0.78	0.69	0.67
ESTOXX		1.00	0.91	0.77	0.83	0.84	0.83
DAX			1.00	0.76	0.80	0.77	0.75
SMI				1.00	0.74	0.66	0.64
FTSE					1.00	0.71	0.69
MIB						1.00	0.76
IBEX							1.00

Table 6: **Equity-Index Returns Correlations.** The correlations are computed from contemporaneous five-minutes returns for the different indices. There is an overlap in trading hours of approximately 9 hours per day, as we exploit the electronic trading in the E-mini contract from 3:15 p.m. to 8:30 a.m. Chicago Time, excluding the 15 minutes trading break between 3:15 p.m. and 3:30 p.m. on Monday-Friday and 4 p.m. to 5 p.m. on Monday-Thursday.

10.5 Parameter Estimates

10.5.1 S&P500

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.963	0.024	c_0^+	4.933	0.064
\bar{v}	0.026	0.000	λ^-	18.266	0.245
κ	8.344	0.135	λ^+	63.193	0.400
σ	0.664	0.018	μ_v	35.125	0.685
κ_u	0.521	0.050	μ_u	60.399	9.101

Panel B: Summary Statistics	
RMSE	2.124
Mean negative jump intensity	1.281
Mean negative jump size	-0.055
Mean positive jump size	0.016
Mean diffusive variance	0.042
Mean negative jump variance	0.008
Mean positive jump variance	0.002

Table 7: **Estimation Results for the Parametric Model defined by Equation (1) for SP Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.2 ESTOXX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.953	0.011	c_0^+	5.077	0.250
\bar{v}	0.033	0.000	λ^-	14.240	0.282
κ	6.689	0.151	λ^+	44.896	0.790
σ	0.689	0.009	μ_v	10.086	0.263
κ_u	1.586	0.078	μ_u	146.338	12.278

Panel B: Summary Statistics	
RMSE	1.897
Mean negative jump intensity	1.729
Mean negative jump size	-0.070
Mean positive jump size	0.022
Mean diffusive variance	0.055
Mean negative jump variance	0.017
Mean positive jump variance	0.005

Table 8: **Estimation Results for the Parametric Model defined by Equation (1) for ES-TOXX Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.3 DAX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.942	0.031	c_0^+	10.291	0.218
\bar{v}	0.031	0.000	λ^-	16.015	0.394
κ	7.050	0.128	λ^+	67.561	0.556
σ	0.659	0.024	μ_v	17.309	0.514
κ_u	1.131	0.085	μ_u	64.924	12.206

Panel B: Summary Statistics	
RMSE	1.859
Mean negative jump intensity	1.464
Mean negative jump size	-0.062
Mean positive jump size	0.015
Mean diffusive variance	0.051
Mean negative jump variance	0.011
Mean positive jump variance	0.005

Table 9: **Estimation Results for the Parametric Model defined by Equation (1) for DAX Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.4 SMI

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.930	0.018	c_0^+	4.117	0.079
\bar{v}	0.022	0.000	λ^-	18.609	0.394
κ	4.915	0.131	λ^+	48.492	0.302
σ	0.468	0.010	μ_v	10.722	0.288
κ_u	1.947	0.109	μ_u	102.066	23.691

Panel B: Summary Statistics	
RMSE	1.618
Mean negative jump intensity	1.963
Mean negative jump size	-0.054
Mean positive jump size	0.021
Mean diffusive variance	0.029
Mean negative jump variance	0.011
Mean positive jump variance	0.004

Table 10: **Estimation Results for the Parametric Model defined by Equation (1) for SMI Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.5 FTSE

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.999	0.016	c_0^+	1.647	0.144
\bar{v}	0.021	0.000	λ^-	15.624	0.401
κ	6.352	0.163	λ^+	37.875	1.401
σ	0.522	0.010	μ_v	28.621	1.218
κ_u	0.911	0.081	μ_u	51.694	16.721

Panel B: Summary Statistics	
RMSE	2.018
Mean negative jump intensity	1.040
Mean negative jump size	-0.064
Mean positive jump size	0.026
Mean diffusive variance	0.042
Mean negative jump variance	0.009
Mean positive jump variance	0.002

Table 11: **Estimation Results for the Parametric Model defined by Equation (1) for FTSE Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.6 MIB

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.969	0.033	c_0^+	3.112	0.072
\bar{v}	0.031	0.000	λ^-	10.059	0.274
κ	4.491	0.139	λ^+	30.659	0.294
σ	0.463	0.015	μ_v	20.840	0.827
κ_u	1.427	0.105	μ_u	5.007	7.572

Panel B: Summary Statistics	
RMSE	2.786
Mean negative jump intensity	0.701
Mean negative jump size	-0.099
Mean positive jump size	0.033
Mean diffusive variance	0.064
Mean negative jump variance	0.014
Mean positive jump variance	0.007

Table 12: **Estimation Results for the Parametric Model defined by Equation (1) for MIB Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.

10.5.7 IBEX

Panel A: Parameter Estimates					
Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.047	c_0^+	6.382	2.349
\bar{v}	0.040	0.001	λ^-	13.603	0.614
κ	3.514	0.121	λ^+	40.044	5.797
σ	0.529	0.027	μ_v	4.397	0.620
κ_u	1.285	0.146	μ_u	55.010	20.512

Panel B: Summary Statistics	
RMSE	2.101
Mean negative jump intensity	2.093
Mean negative jump size	-0.074
Mean positive jump size	0.025
Mean diffusive variance	0.058
Mean negative jump variance	0.023
Mean positive jump variance	0.008

Table 13: **Estimation Results for the Parametric Model defined by Equation (1) for IBEX Options.** **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed 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. Variances and jump intensity are given in annualized decimal units.