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

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

July 12, 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 embeds 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 relative tail risk pricing displays pronounced heterogeneity for the Southern European countries. During the sovereign debt crisis, their stock markets react almost identically to Euro-wide systematic shock, but these events are priced very differently in the respective option markets, indicating differences in crash beliefs across 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.

[†]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. Can they 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 provides an opportunity to shed some light on this question. In particular, the trading of equity-index options has grown sharply in the global financial centers, with both more strikes per maturity and additional maturities on offer. The latest development is a dramatic increase in the trading of options with short tenor. As a result, we now have access to active prices and quotes for financial securities that embed rich information about the pricing of market tail risk in many separate countries. In the current work, we draw on daily observations for option indices in the 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 for the pricing of equity market risk across North America and Europe.

The confluence of turbulent periods renders recent years an excellent "laboratory" for analysis of the way investors treat evolving financial risks, and especially their attitude towards tail risk. Over our sample, three major shocks roiled the global markets, and we exploit the option data to study how tail risk was perceived and priced across these episodes. By combining the pricing of financial risks with ex post information on realized returns, volatilities and jumps, we can gauge the relative size of the risk premiums and what factors drive the risk compensation. The performance of the individual indices varies drastically, with the German market appreciating by an average of about 5% per year and the Italian depreciating by 10% annually. We exploit this heterogeneity to strengthen earlier empirical evidence (Andersen et al. (2015b), Bollerslev et al. (2015)) for the U.S. market regarding the connection between market tail risk and the equity risk premium.

Standard option pricing models capture the dynamics of the equity-index option surface through the evolution of factors that determine the volatility of the underlying stock market, see, e.g., Bates (1996, 2000, 2003), Pan (2002), Eraker (2004) and Broadie et al. (2007). However, recent evidence suggests that the fluctuations in the left tail of the risk-neutral density, extracted from equity-index options, cannot be spanned by regular volatility factors. Hence, a distinct factor is necessary to account for the priced downside risk in the option surface, see Andersen et al. (2015b).[1]

The European samples are shorter and contain fewer options in the strike cross-section, rendering separate day-by-day identification of the factors more challenging than for the S&P 500 index in

---

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

the US. As a consequence, we specify a model with a single volatility component along with the tail factor. This facilitates robust factor identification and reduces the number of model parameters, 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 Andersen et al. (2015b) model closely for US data, 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 remains elevated long after volatility subsides to pre-crisis levels. This feature is in line with Andersen et al. (2015b), but incompatible with the usual approach to the modeling of volatility and jump risk in the literature.

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

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, German volatilities remain 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 on the other, with the latter representing the Southern European countries in our sample.

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

To gauge the risk premium component of the tail risk factor, we complement the option-based evidence with an analysis of high-frequency return data. In particular, we compare the realized jump risks across the euro-denominated indices, DAX, MIB, and IBEX. Our analysis shows remark-

able commonality in the way these markets react to systematic jumps events (when all markets jump simultaneously) over the period 2007-2009, including the global financial crisis. However, in the second part of the sample, involving the sovereign debt crisis, the German index response to systematic jump events is muted relative to those in Italy and Spain. Meanwhile, the latter continue to display near identical responses to such events.

Our analysis shows that the difference 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 outsized 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 equity market return data are not sufficient to capture the crash beliefs of the market participants.

In summary, we find a striking degree of robustness in the pricing of equity 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. 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 readily identifiable.

The rest of the paper is organized as follows. We initiate the analysis with a summary of key empirical results for the US market 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. Section 5 reviews the estimation method, and demonstrates how we extract information and obtain the model-implied factors. In Section 6, we display and discuss 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 the international stock markets. In Section 8, we use the option-extracted factors and the high-frequency futures data to explore the premiums for negative jump risk in the euro-zone countries. Section 9 concludes. Additional details on a variety of aspects related to the data, the estimation and the results are contained in the Appendix.

3

2 Illustrative Evidence for the S&P 500 Index

This section exemplifies our key findings by reviewing results based on S&P 500 index (SPX) options at the Chicago Board Options Exchange (CBOE) and high-frequency data for S&P 500 futures at the CME Group. It extends evidence from Andersen et al. (2015b), but relies on a simplified two-factor model. Moreover, the current 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 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 price jumps and positive volatility jumps coincide. Even if the model is pared down relative to Andersen et al. (2015b), it captures the salient features of the option surface and equity return dynamics across time and markets. Details regarding the data, model and inference techniques are deferred to subsequent sections.

2.1 The U.S. Evidence

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 in the S&P 500 index. 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 current representation operates with a single volatility factor, V, and a separate downside jump intensity factor, U, 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 for the S&P 500 index. The top panel of Figure 1 displays the end-of-day implied spot volatility. The extreme spike in volatility surrounding the financial crisis stands out, while the elevation associated with the intensification of the European sovereign debt crisis in the spring of 2010 and summer of 2011 is 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 am-
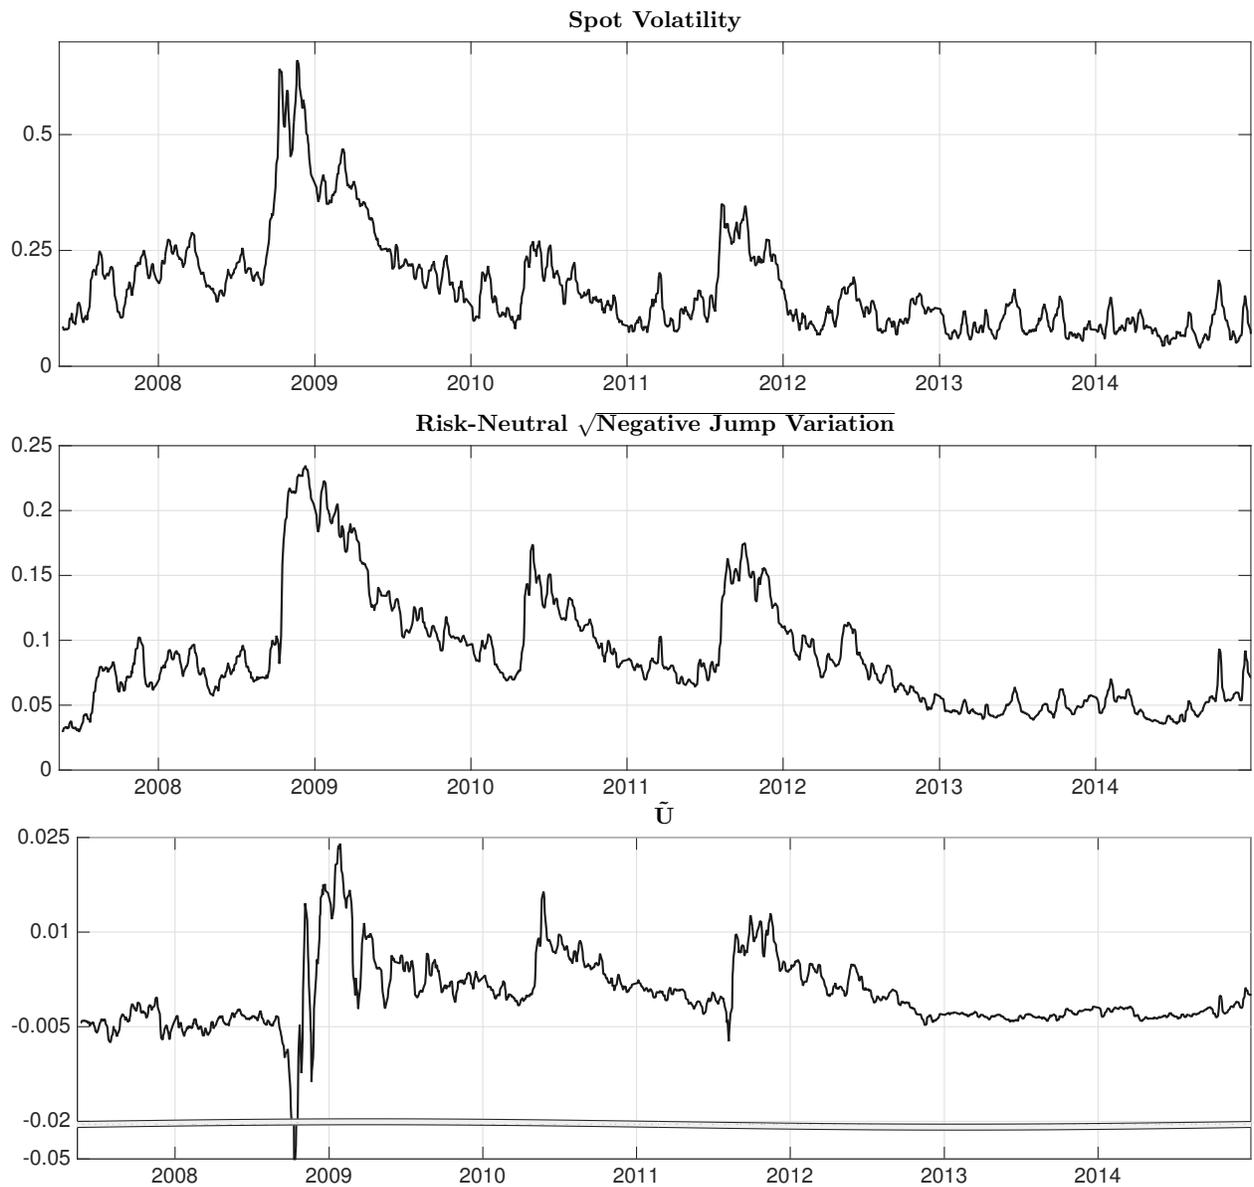
4

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.

plitude.[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 factors 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 $\widetilde{U}$. 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 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 reveals another common trait. Following each major market disruption, the left jump factor is elevated for much longer than spot volatility. In other words, the "excitement" of the left tail lingers well beyond what is observed for volatility. This phenomenon generates 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

---

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

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

power for the state variables, i.e., the implied spot variance, V, and the residual jump intensity, $\widetilde{U}$, 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.
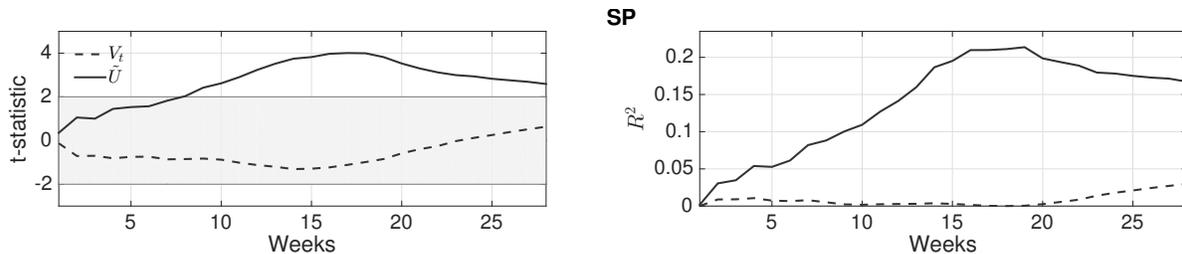
SP

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).

We find the excess returns to be linked to the component of the left jump tail intensity factor unspanned by V, i.e., $\widetilde{U}$, 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 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, $\widetilde{U}$, and the component of the VRP that is orthogonal to the tail factor.

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
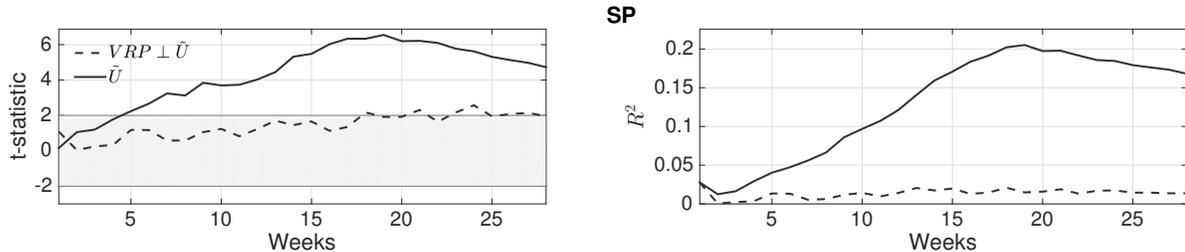
Figure 3: **Predictive Regression for Excess Returns using the Variance Risk Premium.**
The figure reports on regressions of future weekly excess equity-index returns on the option-implied
negative jump intensity (orthogonal to spot variance), $\widetilde{U}$, and the variance risk premium, VRP,
orthogonal to $\widetilde{U}$. The VRP is estimated as the difference between the CBOE VIX squared and the
predicted monthly realized variance (including 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 ($\widetilde{U}$ and VRP). The right panel depicts the R^2
obtained using just the VRP (dashed line) and both $\widetilde{U}$ and VRP (solid line).

the (conditional) equity risk premium.

In contrast, when we regress the future (realized) return variation, stemming from both dif-
fusive volatility and jumps, on the state variables, the relative explanatory power is reversed, as
documented in Figure 4. The jump intensity unspanned by the market volatility has no predictive
power for the return variation, whereas spot volatility is a strong predictor of 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 underlying asset return observations. How-
ever, this factor is unrelated to the equity risk premium which, instead, is tied to the left jump tail
factor. 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 future
market excess returns. Since the 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 identification of the
equity risk premium, corroborating evidence from Andersen et al. (2015b).

These findings raise a number of questions. Are similar option pricing and risk premium pat-
terns 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.
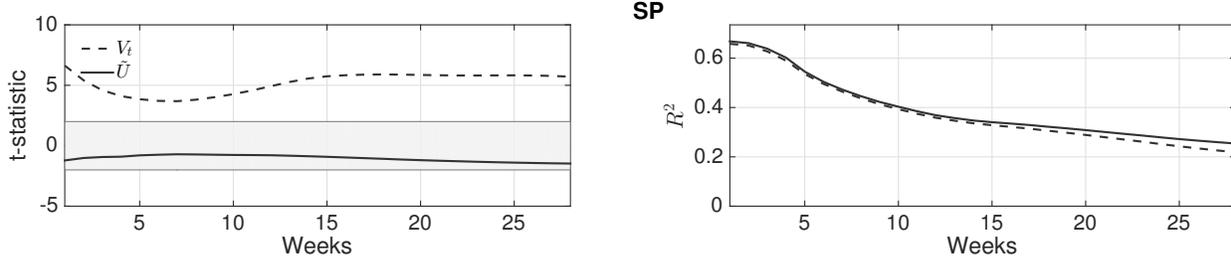
8

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.

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 3 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

9

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 4 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 2 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,

---

[4]The 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.

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 somewhat 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 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.

Further, 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 5 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 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),

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

where $(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}})$ is a two-dimensional Brownian motion with $\mathrm{corr}\left(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}}\right) = \rho$, 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, $\widetilde{\mu}^{\mathbb{Q}}(dt, dx) = \mu(dt, dx) - dt\, \nu_t^{\mathbb{Q}}(dx)$, 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,

---

[5]The 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.

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

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,

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

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), Gruber et al. (2015) 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).[6] For future reference, we label our two-factor affine model, including the negative jump intensity U, the 2FU model.

The 2FU model implies that the risk-neutral instantaneous return variation is governed by the two state variables, V_t and U_t, which control the diffusive and jump variation, respectively. Denoting the (spot) quadratic return variation QV_t, we have,

$$QV_t = V_t + JV_t = V_t + \int_{\mathbb{R}} x^2 \nu_t^{\mathbb{Q}}(dx), \tag{4}$$

The jump variation, JV_t, may be decomposed further into terms stemming from negative and positive jumps, i.e., $JV_t = NJV_t + PJV_t$,

---

[6]There is some option-based evidence, see e.g., Jones (2003) and Christoffersen et al. (2010), that non-affine models work better. Our emphasis here is on short-maturity options and hence potential misspecification stemming from nonlinearities in the dynamics should have relatively small impact on our analysis.

$$NJV_t = \int_{x<0} x^2 \, \nu_t^{\mathbb{Q}}(dx) \qquad \text{and} \qquad PJV_t = \int_{x>0} x^2 \, \nu_t^{\mathbb{Q}}(dx). \tag{5}$$

The intertemporal variation in NJV_t is proportional to U_t, but the level also reflects the jump size distribution. Specifically, for the $2FU$ model, we obtain,

$$NJV_t = \frac{2}{\lambda_-^2} \, U_t. \tag{6}$$

In contrast, the (risk-neutral) positive jump variation, PJV_t, is constant in our model due to the empirically motivated restriction that the positive jump intensity is time-invariant, $c^+(t) = c_0^+$.

Finally, for future reference, we introduce the integrated jump variation for the time interval $[t, t+h], h > 0$, with evident analogies for the negative and positive jump variation measures,

$$JV_{t,t+h} = E_t^{\mathbb{Q}} \left[\int_t^{t+h} \int_{\mathbb{R}} x^2 \, \nu_s^{\mathbb{Q}}(dx) \, ds \right]. \tag{7}$$

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 BSIV is denoted $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$. Estimation of the parameter vector and the period-by-period realization of the state vector now proceeds by minimizing the distance between observed and model-implied BSIV in a metric that also penalizes for discrepancies between the inferred spot volatilities and those estimated (in a model-free way) from high-frequency observations on the underlying asset, $\sqrt{\widehat{V}_t^n}$.[7] The imposition of (statistical) equality between the spot volatility estimated from the actual and risk-neutral measure reflects an underlying no-arbitrage condition which must be satisfied for the option pricing paradigm to be valid.

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

---

[7]We obtain the spot volatility via the so-called truncated realized volatility estimator, using five-minute returns over a three-hour window prior to the close of the trading day. This implementation follows Andersen et al. (2015b), see the Appendix for further details.

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,

$$\left(\{\widehat{\mathbf{S}}_t\}_{t=1}^T, \widehat{\theta}\right) = \operatorname*{argmin}_{\{\mathbf{S}_t\}_{t=1}^T, \theta} \sum_{t=1}^T \left\{ \sum_{\tau_j, k_j} \frac{\left(\kappa_{t,k_j,\tau_j} - \kappa_{k_j,\tau_j}(\mathbf{S}_t, \theta)\right)^2}{N_t} + \frac{\xi_n}{N_t} \frac{\left(\sqrt{\widehat{V}_t^n} - \sqrt{V_t}\right)^2}{\widehat{V}_t^n/2} \right\}, \quad (8)$$

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 (8). 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 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, contains only a single volatility factor. One potential concern is that the relatively constrained modeling of the volatility process

15

may induce a bias in the extraction of the implied jump dynamics from the option surface. For the S&P 500 sample, we have a natural benchmark. Andersen et al. (2015b) estimate an extended version of model (1) which provides an excellent fit to the option prices as well as the volatility series, clearly outperforming more standard and equally heavily parameterized representations.

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

Because the left tail factor enters the specification for the jump intensity in distinct ways, a direct comparison of the factor realizations across the two models is not meaningful. Instead, we focus on the model-implied variation in the negative jump intensity stemming from the left tail factor. In both cases, the effect is proportional to the concurrent value of U factor in the model. Consequently, Figure 5 depicts the standardized jump factor series extracted from the two separate models, estimated from partially overlapping periods and option samples. The models deliver remarkably similar time series paths for the return variation induced by the left tail factor, with a correlation of 0.987 during the overlap, corroborating the robust identification of the jump tail factor from either model.

6.2   Country-by-Country Factor Realizations

We now provide more details regarding the implied factors extracted from the equity-index option markets based on our $2FU$ model (1). We first compare the implied spot variance for 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 volatility factors from the Euro ESTOXX (Eurozone), DAX (Germany), SMI (Switzerland), FTSE 100 (U.K.), FTSE MIB (Italy), and IBEX 35 (Spain) markets. The most striking feature is the extraordinary close association between many of these factors and the S&P 500 spot volatility, not just in terms of correlation, but also level. For example, the U.K. volatility is barely distinguishable from the S&P 500 factor throughout the sample, while the Swiss factor deviates only for a few episodes following the constrained Swiss franc-euro exchange rate policy implemented in September 2011. In the former case, the volatility correlation is about 98% and in the latter around 95%. In contrast, the discrepancies between the
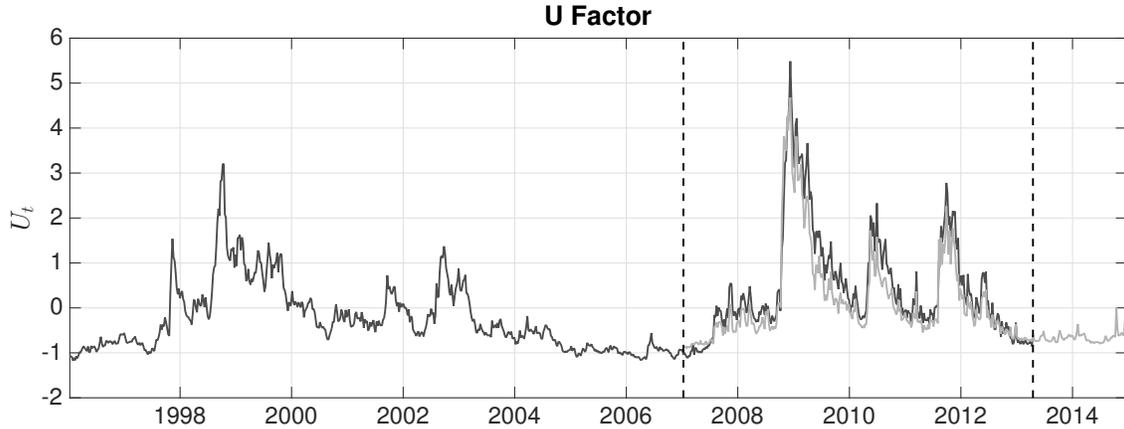
Figure 5: **Comparison of Model-Implied Left Jump Tail Factors.** The figure depicts the daily standardized option-implied left jump tail factors from model (1) and the Andersen et al. (2015b) model. The series are extracted based on parameter estimates from weekly SPX option prices observed over January 2007–December 2014 and January 1996-July 21, 2010, respectively. Both series are normalized to have mean zero and unit variance.

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 from that point onward, albeit to a varying degree. For the broader euro-zone ESTOXX index, the same effect is clearly visible and originates with the 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 heterogeneous. In addition, the relative impact across the 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 these two Southern European indices attain a plateau well above the others ever since the sovereign debt problems surfaced in early 2010. This systematic divergence over the second half of our sample lowers the volatility correlations for MIB and IBEX with S&P 500 to 0.75 and 0.77, respectively.

Next, Figure 7 depicts the option-implied (risk-neutral) negative jump variation, NJV_t, for these indices along with the corresponding quantity for the S&P 500. As noted in Section 4, the risk-neutral negative jump variation equals $\frac{2}{\lambda_-^2}U_t$ in the $2FU$ model. Hence, these jump variation series are directly proportional to the negative jump intensity factor, U. Consequently, the relative variation in the series directly reflects the corresponding variation in the extracted jump intensity state variables for the individual indices.

17
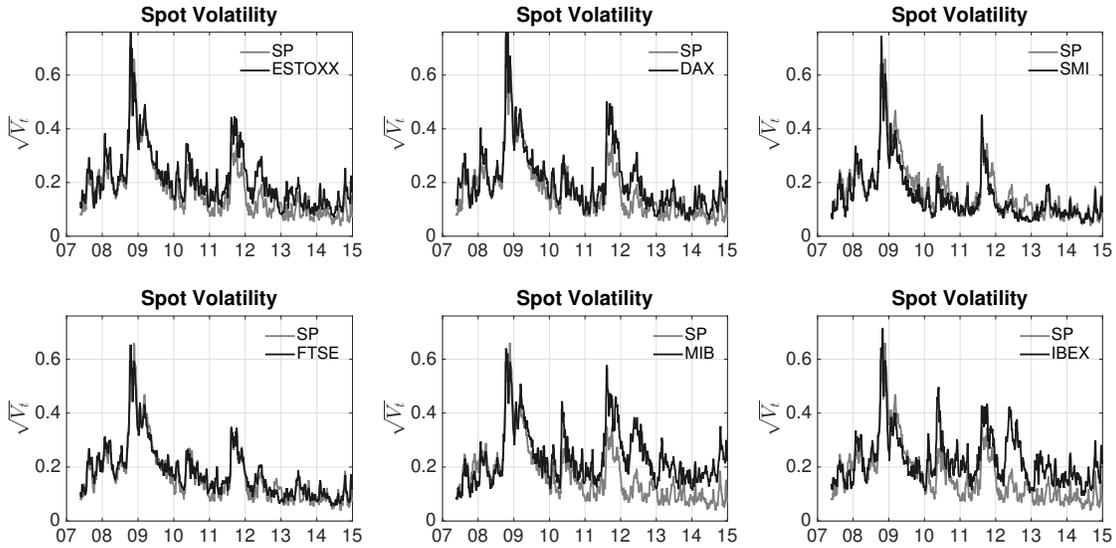
Figure 6: **Volatility Factor Comparison.** For each option-implied spot variance daily factor, we report the trailing five-day moving average of $\sqrt{V_t}$. The pairwise correlation between the volatility factor of the S&P 500 and each of the European indices is as follows: SP-ESTOXX: 0.94; SP-DAX: 0.93; SP-SMI: 0.94; SP-FTSE: 0.98; SP-MIB: 0.77; and SP-IBEX: 0.78.

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 $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx) \equiv \frac{2}{\lambda_-^2} U_t$. The pairwise correlation between the jump factor of the S&P 500 and each of the European indices is as follows: SP-ESTOXX: 0.96; SP-DAX: 0.97; SP-SMI: 0.95; SP-FTSE: 0.98; SP-MIB: 0.84; and SP-IBEX: 0.78.

Qualitatively, the pattern is similar to the one observed for the volatility factors. For each index, the volatility and jump 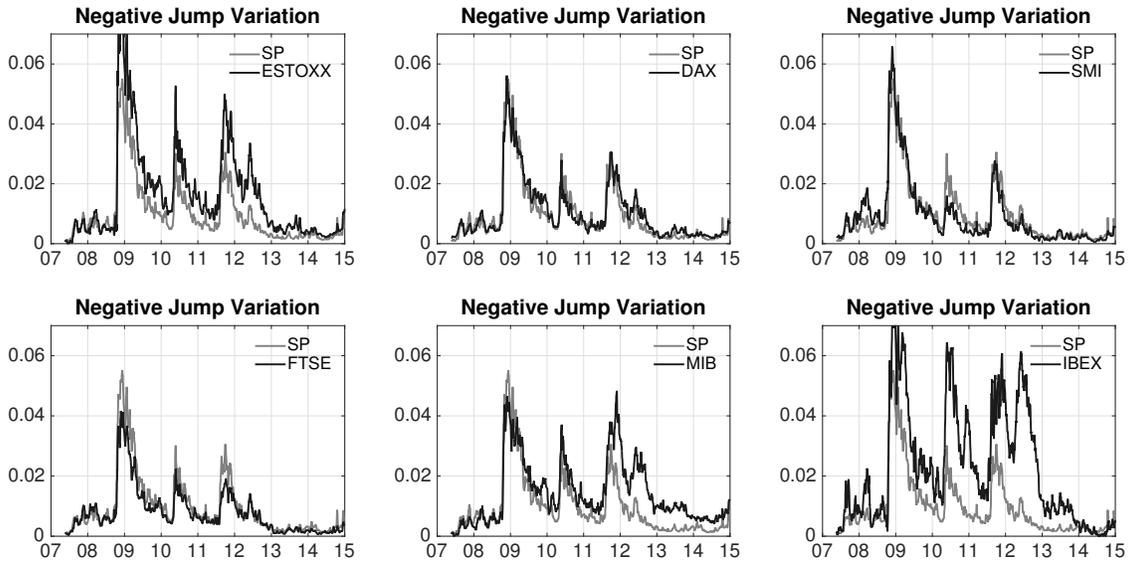factors are highly correlated. Nevertheless, as for the U.S., the relative size of the spikes in the jump intensity versus volatility varies substantially, with the European crisis inducing a stronger surge in the jump intensity relative to (diffusive) volatility. Again, the U.K. and U.S. series evolve in near unison and display a correlation above 98%, even if the British jump variation is slightly lower throughout. Similarly, the jump intensity factor for the Swiss index correlates strongly with the S&P 500 factor, although the Swiss factor tends to be above the one in the U.S. 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 add further credence to the robustness of our methodology in extracting the salient pricing features from the option surfaces.

7   Option-Based Prediction of the Future Risk and Risk Premium

We now explore the ability of the option-implied factors, the spot variance and left jump intensity, to forecast the (realized) return variation – defined as the sum of the squared high-frequency equity-index futures returns – and the equity excess return. The former signifies whether the factors are associated with market-wide risk as reflected by the realized equity volatility and jump activity, while the latter speaks to their forecast power for future equity returns. Since many recent studies identify the variance risk premium as a key explanatory variable for the equity risk premium, we also compare the predictive performance of our jump factor to that of the variance risk premium.

We rely on predictive regressions to assess the forecast power of the state variables. As for the S&P 500 index, we mitigate the look-ahead bias by re-estimating the parameters of the $2FU$ 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.[8] Specifically, we extract the state vector realizations from the objective function (8) by optimizing over $\mathbf{S}_t$ for each date t over 2008-2014, with θ fixed at the estimate obtained from the 2007 data.

7.1 Predicting the Equity Risk Premium

This section focuses on the evidence for predictability of the equity risk premium. We regress the future excess returns for each equity index on their option-implied state variables. In all instances, the two state variables are correlated so, for ease of interpretation, we supplement the spot variance factor, V, with the component of the left jump factor that is orthogonal to the spot variance, denoted $\tilde{U}$, as the second explanatory variable. This approach ascribes 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 $t = 1, \ldots, T - h$,

$$r_{t,t+h} \;=\; \log(X_{t+h}) - \log(X_t) \;=\; c_{0,h} \;+\; c_{v,h} \cdot V_t \;+\; c_{u,h} \cdot \tilde{U}_t \;+\; \epsilon_{t,h}\,, \tag{9}$$

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 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

---

[8]The 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.

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 multi-horizon excess returns and return variation measures. The inference is based on Newey-West robust standard errors, accounting for the estimation errors in any potential first-stage regression used to normalize the regressors.

The regression results are qualitatively similar across all indices. For example, the top left panel of Figure 8 plots the t-statistics for the ESTOXX regression slopes in equation (9), while the top right panel displays the corresponding R^2. Consistent with the hypothesis of a slowly evolving equity risk premium, the predictive power rises steadily with the horizon. The unpredictable return component seems to dominate at the weekly horizons, but as the noise innovations diversify over time, the predictable component emerges over the longer holding periods. 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, as 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. This conclusion is 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 the board. Apart from the 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.

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 (9), but the dependent variable is now a measure of the future realized return variation over the time interval $[t, t + h]$,
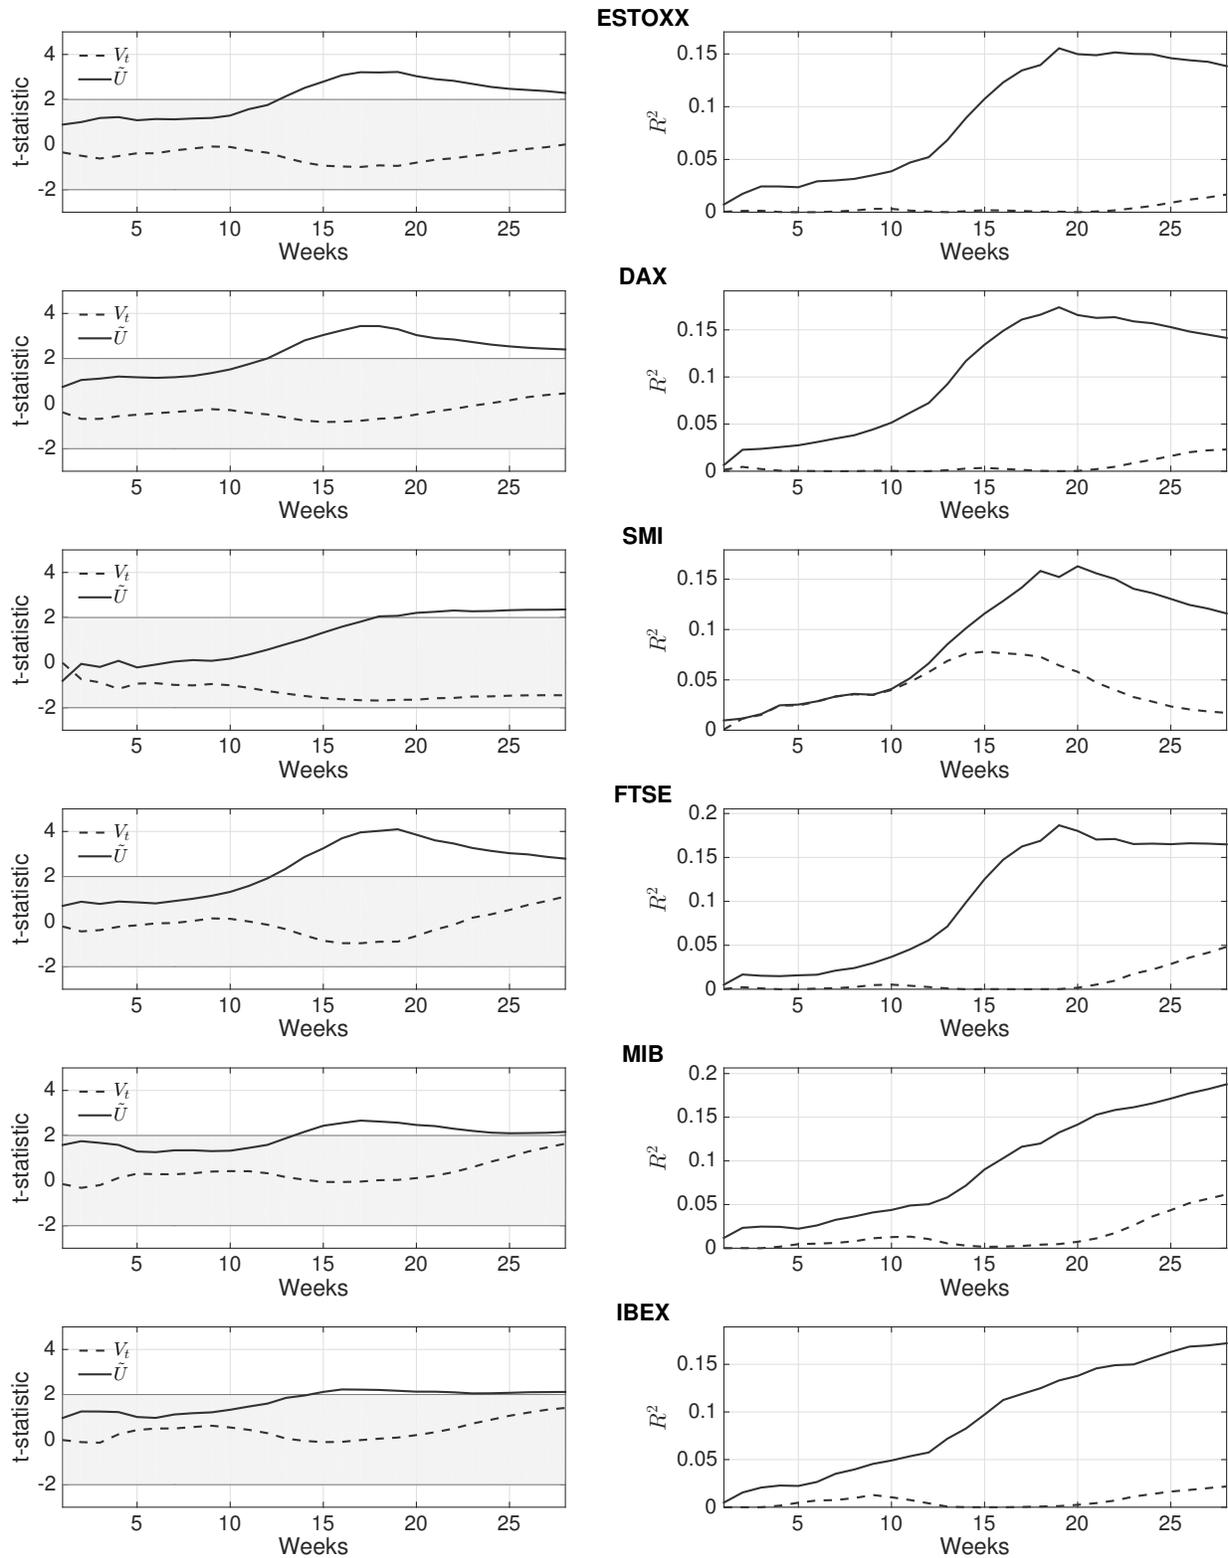
Figure 8: **Predictive Regressions for Excess Returns using the State Variables**. Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2, where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

$RV_{t,t+h}$. 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,

$$RV_{t,t+h} \; = \; k_{0,h} \, + \, k_{v,h} \cdot V_t \, + \, k_{u,h} \cdot \widetilde{U}_t \, + \, \epsilon_{t,h} \, , \tag{10}$$

The left panels of Figure 9 reveal, in line with the prior US 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.[9] The variance risk premium is the compensation demanded by investors for bearing variance risk.[10] Formally, it is defined as the gap between the conditional risk-neutral and statistical expectation of the 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 the value of the VIX index.[11] 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, in light of our findings above,

---

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

[10]See also Bakshi et al. (2011) for the predictive ability of other option-based volatility measures.

[11]See for example Bakshi and Madan (2000) and Carr and Wu (2009) for the construction of model-free measures from options portfolios.
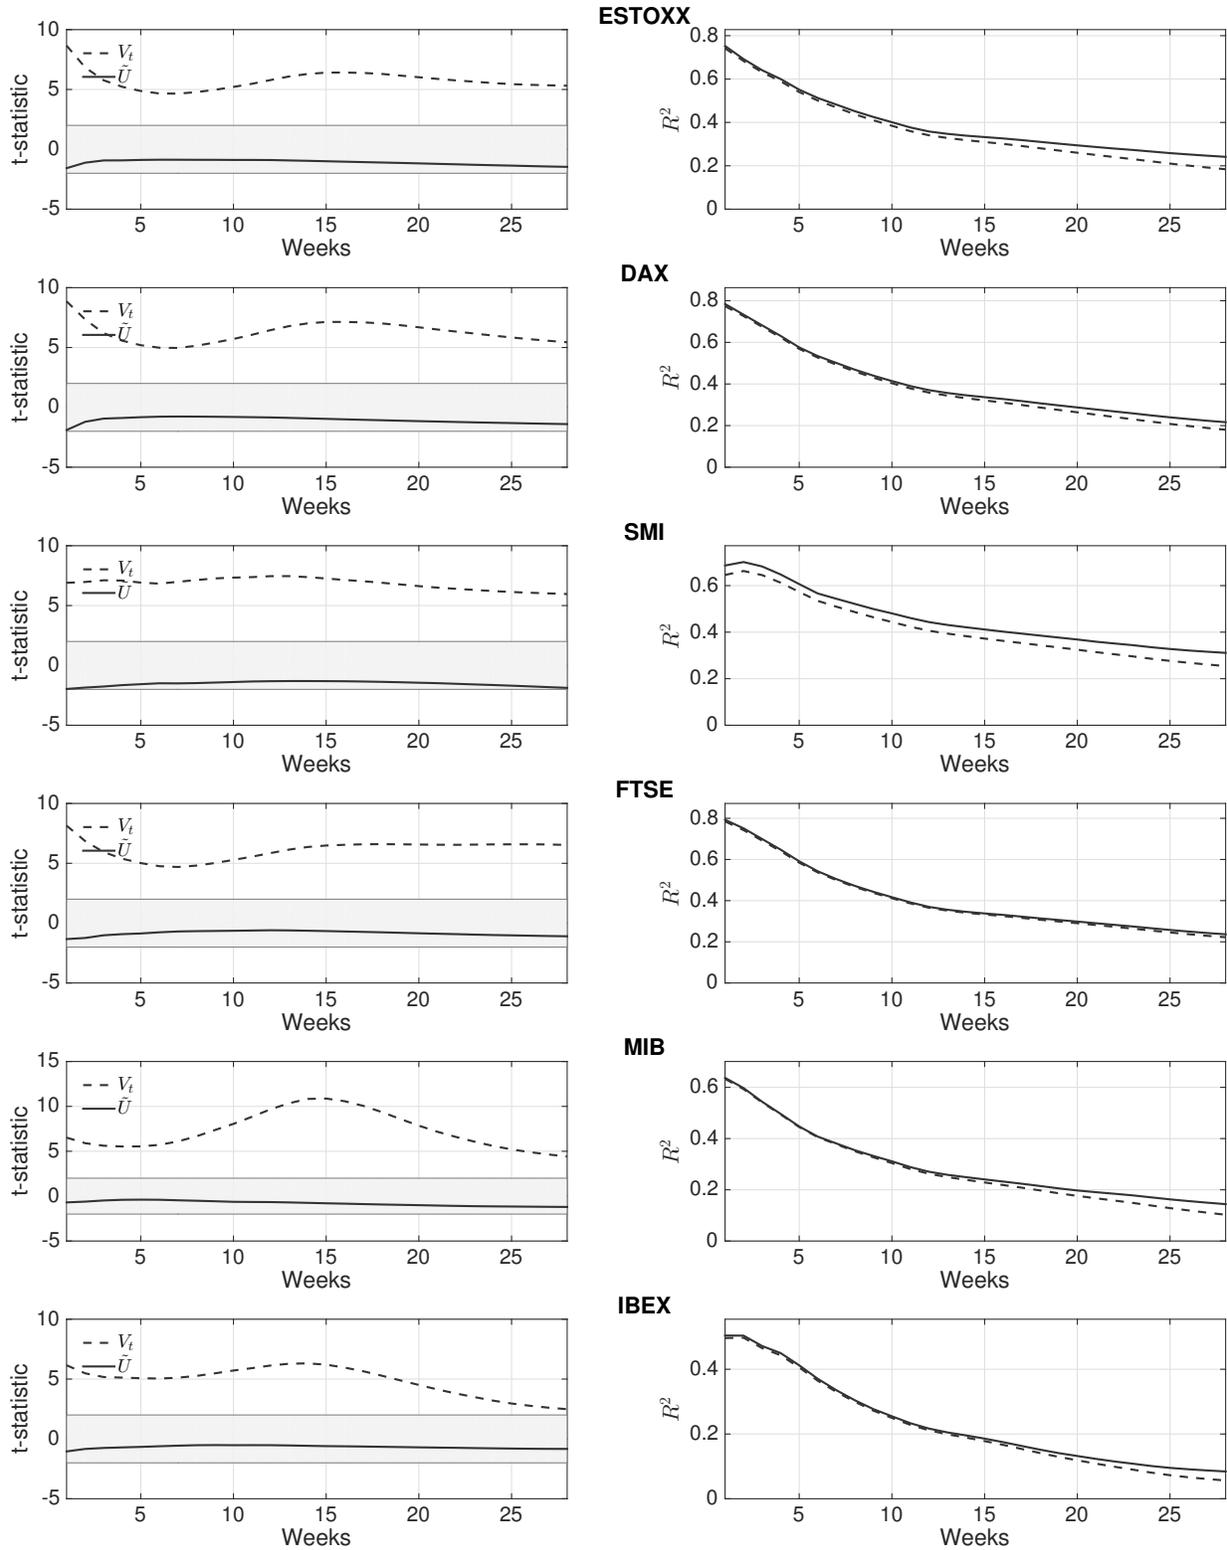
Figure 9: **Predictive Regressions for Return Variation**. Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2, where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

it is natural to ask whether the predictive ability of the variance risk premium is due solely to the
"pure tail" factor, i.e., $\widetilde{U}$.

To address this question, we 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).[12]

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,

$$ r_{t,t+h} \;=\; c_{0,h} + c_{u,h} \cdot \widetilde{U}_t + c_{p,h} \cdot \widetilde{VRP}_t + \epsilon_{t,h} \,, \tag{11} $$

where $\widetilde{VRP}_t$ denotes the (estimated) variance risk premium at the end of trading day t, orthogo-
nalized 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, if anything, lower than before and the
jump factor remains highly significant for the longer return horizons. Finally, there is no systematic
evidence of auxiliary predictive power of the variance risk premium over-and-above the jump factor.
We conclude that the jump factor is the source of the predictive power of the variance risk premium
for future equity returns. We strongly reject the reverse hypothesis that the variance risk premium
is a superior predictor relative to our jump intensity factor.[13]

8 Downside Jump Risk and Risk Premium Linkages

We have identified the jump factor as a primary indicator of the future equity risk premium. This
factor represents the risk-neutral negative jump intensity – or more generally the negative jump
variation – so any increase may stem from two separate sources. It may be due to an elevation in
the expected negative jump variation under the statistical (objective) measure ($\mathbb{P}$) or, alternatively,

---

[12]Specifically, 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 $-8 \le m \le 8$ with
increments of 0.1. We then interpolate the implied volatility as a function of K. When K is lower (higher) than
the lowest (highest) available strike, we extrapolate the implied volatility outside the defined moneyness range as a
constant equal to the implied volatility at the lowest (highest) available strike. We obtain the VIX index for both
maturities and linearly interpolate to obtain the VIX index corresponding to 30 calendar days.

[13]If we include the VRP and the pure tail factor orthogonalized with respect to VRP as explanatory variables, the
tail factor has significant auxiliary predictive power relative to VRP for the equity returns at the 3-6 month horizon.
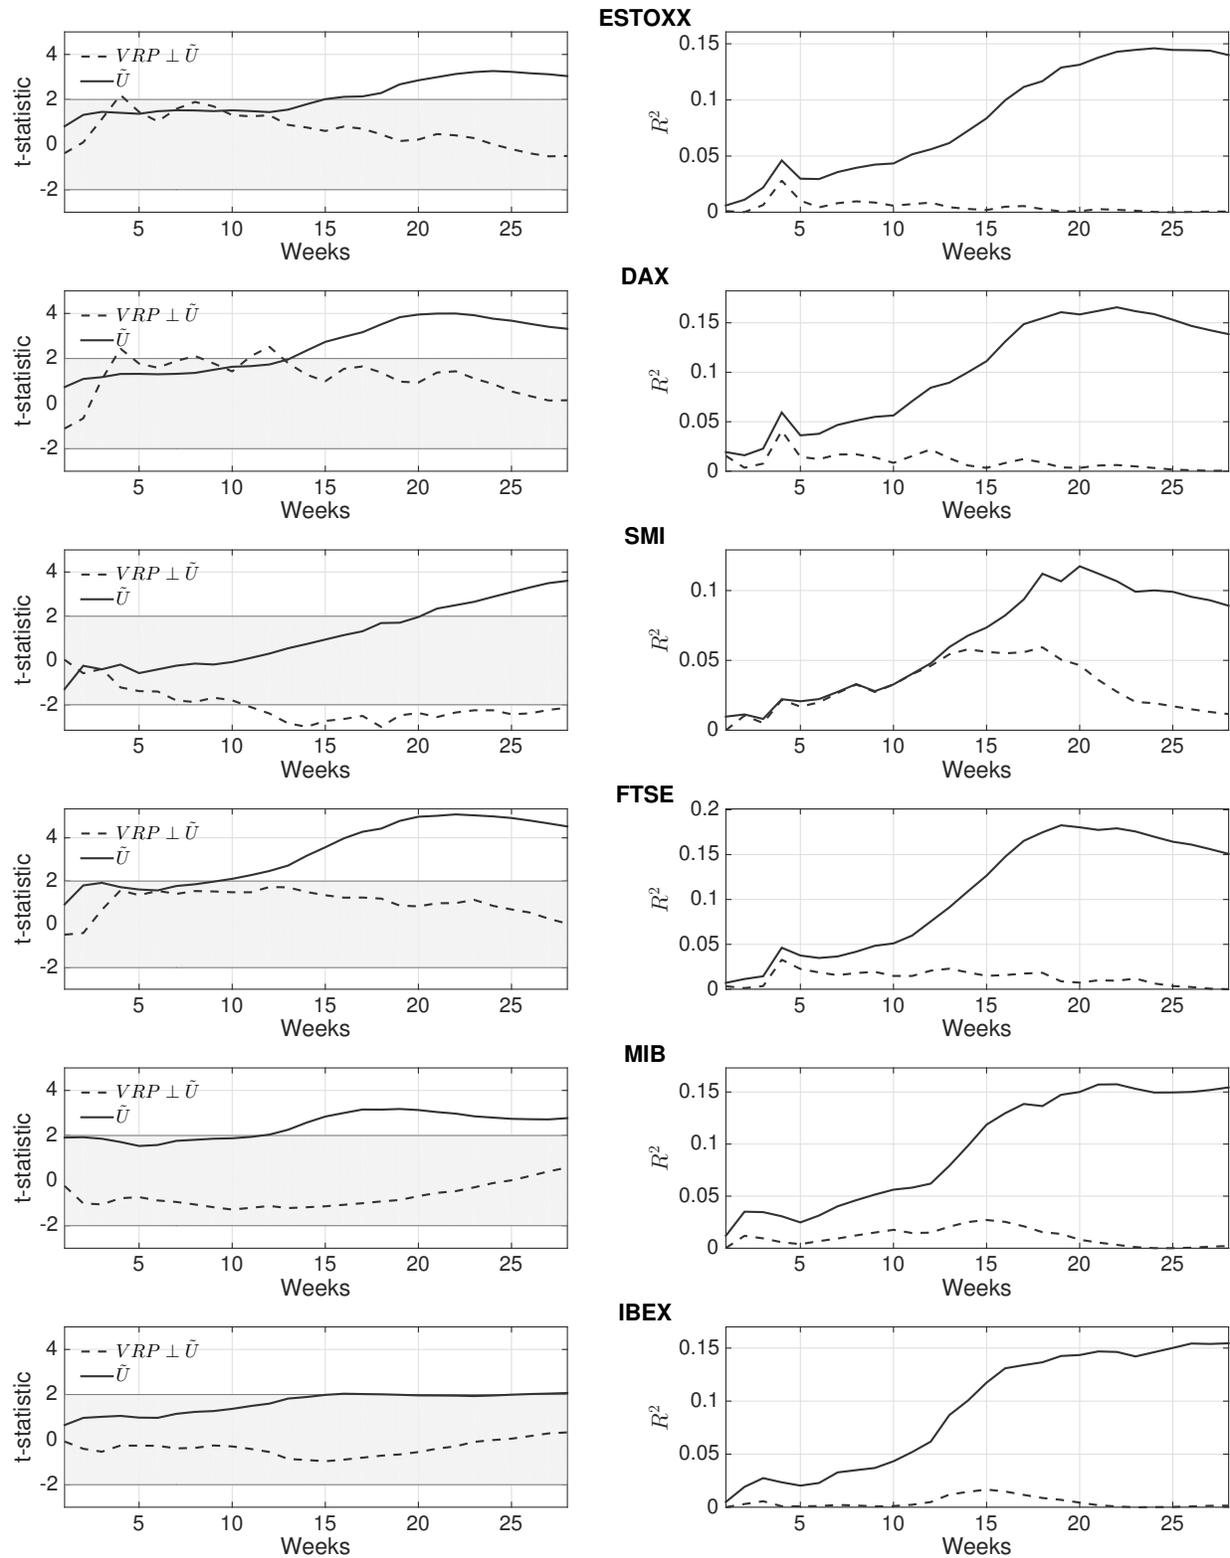
Figure 10: **Predictive Regressions of Excess Returns using the Variance Risk Premium**. Left Panel: t-statistics for the regression slopes; Right Panel: Regression R^2, where the full drawn line depicts the total degree of explained variation and the dashed line represents the auxiliary explanatory power afforded by the VRP.

it may reflect a rise in the compensation for downside tail risk, i.e., the gap between the negative jump variation under the $\mathbb{Q}$ (risk-neutral) and $\mathbb{P}$ measure.

From an economic perspective, it is of interest to disentangle these effects. Unfortunately, while risk-neutral expectations of future jump risk are easy to generate – and readily provided by our parameter and factor estimates – it is difficult to compute reliable estimates for their counterparts under the statistical measure, 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 conditional price dynamics under the objective measure, so those parametric assumptions are much harder to assess on a day-to-day basis. This raises serious concerns regarding model misspecification, adding an element of ambiguity to the interpretation of the empirical findings.

In view of the above difficulties, we turn to nonparametric estimators for the jump variation measures and, acknowledging that they provide noisy estimates for the *expected* jump variation, we abstain from any attempt at estimating the *level* of the negative jump variation and the associated premium for the individual indices. Instead, we may infer a great deal about the *relative* exposure of the indices to shocks that induce actual jumps in the stock markets. This may be achieved, throughout this eventful period for the global markets, by measuring whether they display similar or divergent jump behavior. To the extent their responses are well aligned, it is natural to infer that the expected jump variation measures were also highly correlated.

Corresponding to the notions of jump variation introduced in Section 4, we define, e.g.,

$$RNJV_{t,t+h} = \frac{1}{h} \left(\sum_{s \in [t,t+h]} (\Delta X_s)^2 1_{\{\Delta X_s < 0\}} \right). \tag{12}$$

This realized negative jump variation measure provides a noisy estimator of the, a priori, expected negative jump variation. If it is highly correlated with the corresponding measure for another index, their ratio is indicative of the relative exposure to shocks triggering jumps.[14]

Nonetheless, there are complications. First, in Section 6.2, we document a strong qualitative correspondence of the negative jump factor across all indices during 2007-2009. However, thereafter,

---

[14]We note that $RNJV_{t,t+h} - NJV_{t,t+h}$ provides a noisy measure of the integrated negative jump variation premium.

in 2010-2014, the Italian and Spanish indices display greatly elevated jump intensities in the periods surrounding the sovereign debt crisis. Most strikingly, the Spanish factor is higher, but roughly proportional with, the other jump factors through 2009. In contrast, after the onset of the debt crisis, the Spanish jump factor fluctuates violently and frequently attains truly extreme levels. This suggests the presence of a structural break in 2010, with the jump exposures diverging significantly from the pattern observed over the first three years of the sample.

A second complication is the difference in currency denomination. 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 the German (DAX), the Italian (MIB) and the Spanish (IBEX) indices are denominated in the identical (euro) currency. These indices further display a divergent dynamic in the pricing of jumps across the sample, rendering them excellent candidates for an exploration of shifts in the relative pricing of jump risk over time.

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

This section focuses on the three euro-denominated indices and explores the first and second part of the sample separately to allow for a potential break in the jump exposure of the indices.

To assess the relative exposure to the jump factor, we extract jump series for each market, using five-minute index futures and the jump identification techniques detailed in the Appendix. We restrict the sample to common exchange trading hours, so that we can monitor the indices simultaneously. Specifically, we classify the jumps for a given index as common or country-specific depending on whether *both* of the other indices jump within the same five-minute interval or not. The common jumps enable us to gauge the relative sensitivity to shocks that impact all indices, while the country-specific jumps reflect more idiosyncratic exposure.

Figure 11 provides pairwise scatter plots of the 344 common jumps for the three indices over the period 2007-2009, encompassing the great recession, but not the sovereign debt crisis. It reveals a remarkable similarity in jump responses. All regression lines have a slope close to, and insignificantly different from, unity suggesting they move one-for-one. The associated R^2 values range from 0.88 to 0.94. In brief, the common jumps reflect near identical index exposures to events with euro-wide implications in the period surrounding the financial crisis.

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 monthly country-specific realized negative and positive jump variation. These measures are depicted in Figure 12. Focusing on the 2007-2009 period, the most noteworthy discrepancy is the limited
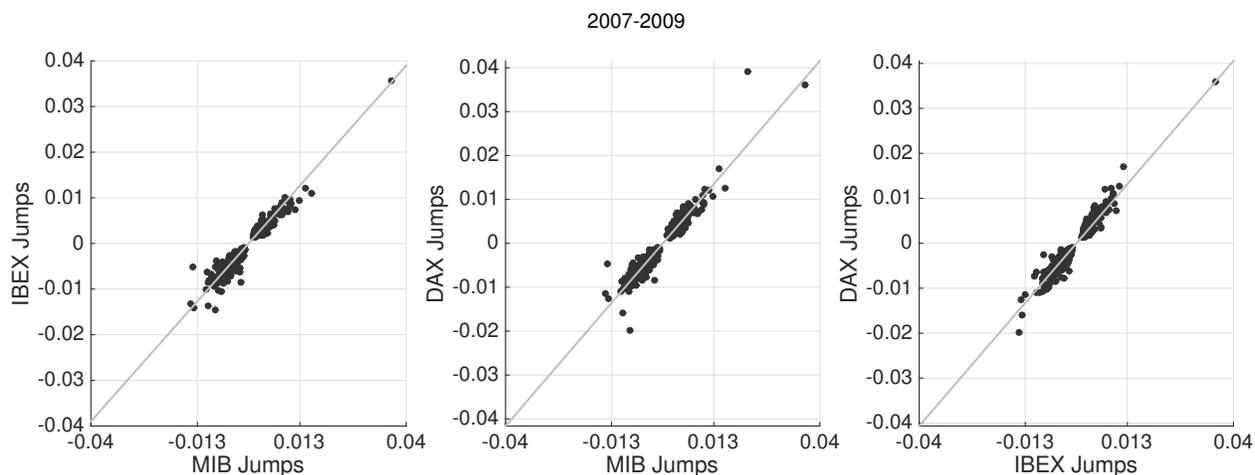
Figure 11: **Scatter of Common Market Jumps in Germany, Italy and Spain**. The straight lines on the plots represent the fit from a linear regression (without an intercept) of the jumps on the Y-axis on the jumps on the X-axis.

degree of positive jump variation in the DAX index. However, on the downside, the DAX jump variation spikes along with IBEX at the beginning of 2008 and during the depth of the financial crises, while the MIB response is smaller in the first instance but more dramatic in the second. Finally, after October 2008, the DAX index appears less susceptible to large negative idiosyncratic shocks than the other two. Nonetheless, overall, for the period from 2007 through the main part of the financial crisis, there is no evidence of significant discrepancies in the realized negative jump exposures of these indices. This is corroborated by the summary statistics in Table 2. In particular, it is evident that the correlation of the DAX measures with the other two are stronger for the negative than the positive jump variation.[15]

Figure 13 provides the scatter plots for the 547 common jumps over 2010-1014. Again, the jump realizations are approximately linearly aligned. However, the DAX responses are now more muted than for the Southern European counterparts. The slopes for the middle and right panels 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 relative German exposure has declined sharply. This is consistent with the sovereign debt crisis having a disproportional impact on Italy and Spain. Moreover, the R^2 values for the DAX plots, at 0.87 and 0.84, have dropped relative to before, while the value for IBEX-MIB is 0.94 and, if anything, higher than during the financial crisis, speaking to a great deal of commonality in the

---

[15]This is consistent with evidence from lower frequency data, e.g., Longin and Solnik (2001), Ang and Chen (2002), Poon et al. (2004), Bae et al. (2003) and recent evidence from high-frequency stock data, e.g., Bollerslev et al. (2013).

Country-specific Jump Variation
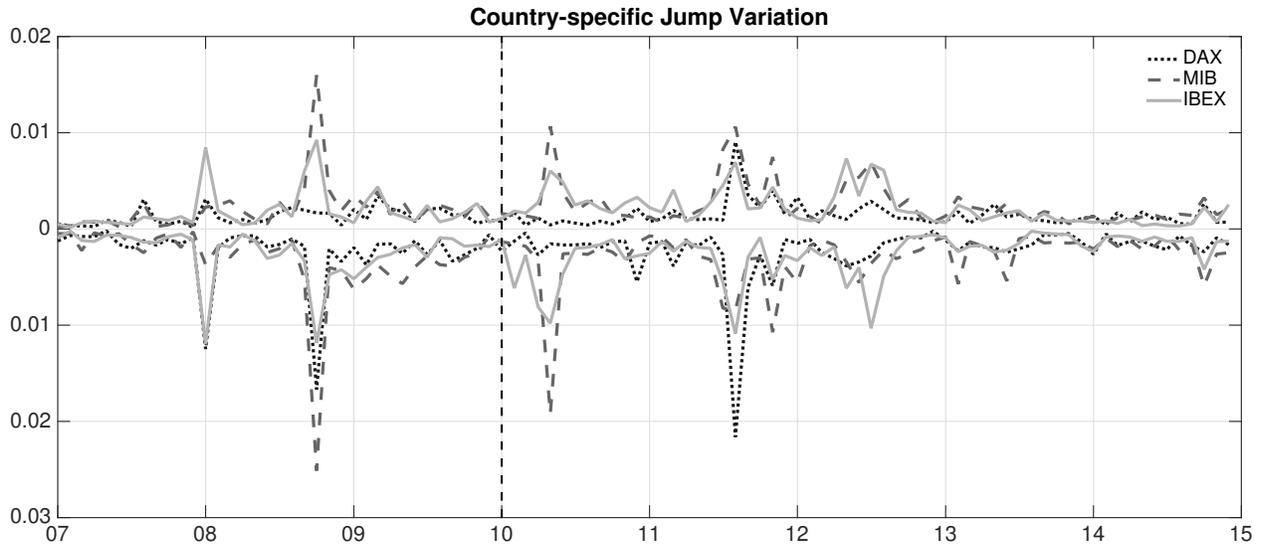
Figure 12: **Country-specific Jump Variation in Germany, Italy and Spain**. The jump variation due to the positive and negative jumps is above, respectively, below the zero line. The jump variations are computed over one month and reported in annualized decimal units. The dashed vertical line separates the 2007-2009 and 2010-2014 sample periods.

jump exposures for Italy and Spain during the debt crisis.

2010-2014
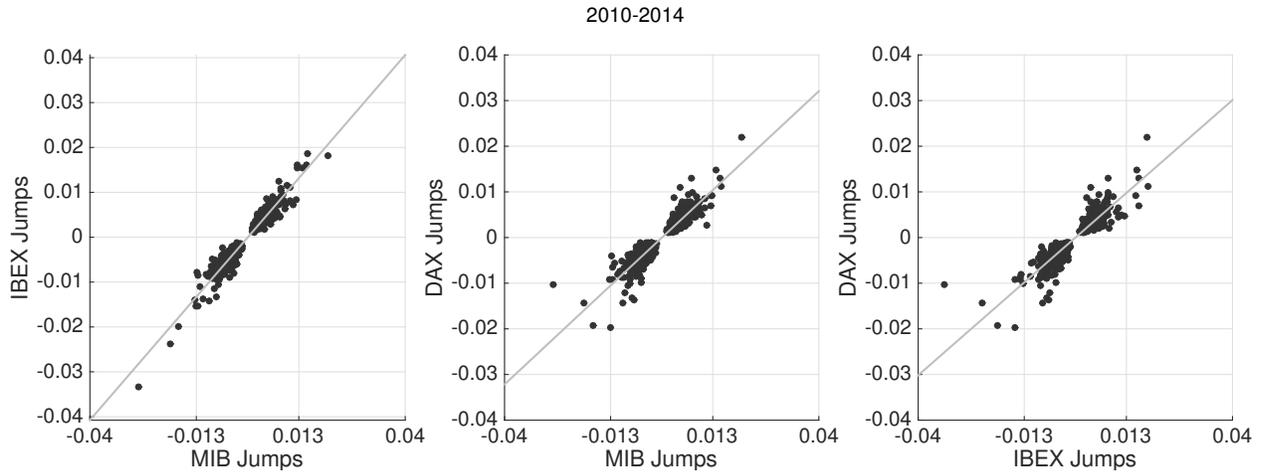
Figure 13: **Scatter of Common Market Jumps in Germany, Italy and Spain**. The straight lines on the plots represent the fit from a linear regression (without an intercept) of the jumps on the Y-axis on the jumps on the X-axis.

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

30

jump exposure is now 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 realized German jump variation is almost uniformly the lowest. 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 shocks than the other. These conclusions are confirmed by Table 2. 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.[16]

## 8.2   Implications for European Jump Risk Premiums

Our nonparametric analysis of the jump risks across Germany, Italy and Spain reveal strong linkages, but also a distinct shift in relative exposures following the financial crisis. By combining the 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. To facilitate this analysis, Figure 14 displays the difference between the option-implied negative jump variation and volatility for the three indices.

It is evident from Figures 7 and 14 that, over the first part of our sample, the risk-neutral negative jump variation of the Spanish index is more elevated and displays more abrupt fluctuations than the Italian, German and U.S. ones. This contrasts sharply 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 that the corresponding premium in Germany and Italy. One may conjecture that the rising tension in housing and mortgage markets, preceding the financial crisis, were taking a disproportionate toll on the risk attitude towards the Spanish economy which was fueled by a real estate boom. In contrast, if anything, the Italian risk premium for negative jump risk is 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. Similarly, the gap between the DAX and MIB jump factor is mostly positive until

---

[16]The 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		
2007-2009					
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
2010-2014					
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 2: **Summary Statistics for the Country-Specific Jump Variation**. The mean and standard deviation are computed from monthly realized measures and reported in annualized decimal units. The top panel covers January 2007 - December 2009, and the bottom panel refers to January 2010 - December 2014.

2010 in Figure 14. This likely reflects a perception of a relatively lower vulnerability of the Italian economy to the sharp contraction in the financial sector.
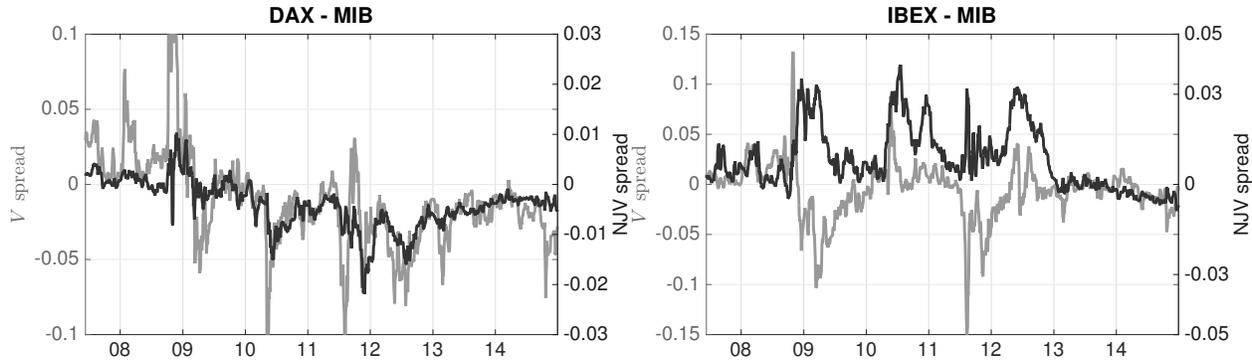
Figure 14: **Option-Implied Factor Spreads in the Euro-Zone**. The left panel depicts the difference between the implied factors (spot volatility and left jump intensity) for the German DAX and Italian MIB index extracted from model (1). The right panel displays the corresponding extracted factor differential between the Spanish IBEX and Italian MIB index.

After 2010, the picture shifts dramatically. Figure 7 shows 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. Figure 14 corroborates this finding and further demonstrates that the jump and volatility factor differentials move largely in tandem, with both generally being smaller in Italy than Germany prior to 2010 and larger thereafter.

As noted previously, the shifts in the Spanish jump factor are even more dramatic. Given the near identical realized negative jump variation in Spain and Italy across the sample, documented in Figure 12–13 and Table 2, it is striking to observe in Figure 14 that, from the onset of the financial crisis in 2008 until the European crisis abates in 2013, the Spanish jump factor exceeds the one in Italy by a substantial margin. We conclude that the Spanish premium for negative jump risk was considerably higher over this period. In contrast, 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 than the other indices during the European debt crisis.

Finally, we note from Figure 14 that the Italian volatility factor exceeds those in Germany and Spain over 2010-2013. It suggests a consensus about persistent economic distress in Italy, even as concerns (and the premiums) for large negative shocks to the Spanish market were more pronounced. One explanation is the amplification of the impact of the sovereign debt crisis on Italy given its large domestic debt burden, while the threat of abrupt drops in the IBEX may be associated with the possibility of Spain dispensing with the euro currency and joining a set of other nations in adopting a Southern euro in its place. Such scenarios were deemed less plausible in the case of Italy, a founding member of the initial treaties seeking increased economic and political

33

integration across Europe.

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.

Bakshi, G. and D. Madan (2000). Spanning and Derivative Security Valuation. *Journal of Financial Economics 55*, 205–238.

Bakshi, G., G. Panayotov, and G. Skoulakis (2011). Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios. *Journal of Financial Economics 100*, 475–495.

Bates, D. S. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies 9*, 69–107.

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

Bates, D. S. (2003). Empirical Option Pricing: A Retrospection. *Journal of Econometrics 116*, 387–404.

Bekaert, G. and R. Hodrick (1992). Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets. *Journal of Finance 47*, 467–509.

Bollerslev, T., J. Marrone, L. Xu, and H. Zhou (2014). Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence. *Journal of Financial and Quantitative Analysis 49*, 633–661.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies 22*, 4463–4492.

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

Bollerslev, T., V. Todorov, and 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 K. Mimouni (2010). Models for S&P 500 Volatility Dynamics: Evidence from Realized Volatility, Daily Returns, and Option Returns. *Review of Financial Studies 23*, 3141–3189.

Christoffersen, P., K. Jacobs, and C. Ornthanalai (2012). Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options. *Journal of Financial Economics 106*, 447–472.

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics 7*, 174–196.

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

Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance 59*, 1367–1403.

Ferson, W. and C. Harvey (1993). The Risk and Predictability of International Equity Returns. *Review of Financial Studies 6*, 527–566.

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

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

Gruber, P., C. Tebaldi, and F. Trojani (2015). The Price of the Smile and Variance Risk Premia. Working paper, University of Lugano and University of Geneva.

Harvey, C. (1991). The World Price of Covariance Risk. *Journal of Finance 46*, 111–157.

Hjalmarsson, E. (2010). Predicting Global Stock Returns. *Journal of Financial and Quantitative Analysis 45*, 49–80.

Jones, C. (2003). The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets. *Journal of Econometrics 116*, 181–224.

Kou, S. (2002). A Jump- Diffusion Model for Option Pricing. *Management Science 48*, 1086–1101.

Li, J. and G. Zinna (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 3: **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 4: **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 $0, \frac{1}{n}, \frac{2}{n}, \ldots$. 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 (8), 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 5: **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	-1.000	0.012	c_0^+	7.551	0.622
$\overline{v}$	0.029	0.000	λ^-	13.548	0.196
κ	6.216	0.086	λ^+	78.255	2.368
σ	0.605	0.008	μ_v	22.663	0.462
κ_u	0.600	0.062	μ_u	18.656	5.389

Panel B: Summary Statistics

RMSE	1.909
Mean negative jump intensity	0.853
Mean negative jump size	-0.074
Mean positive jump size	0.013
Mean diffusive variance	0.038
Mean negative jump variance	0.009
Mean positive jump variance	0.002

Table 6: **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 statisticsfor 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	-0.999	0.014	c_0^+	12.879	0.777
$\overline{v}$	0.032	0.000	λ^-	14.569	0.180
κ	7.471	0.170	λ^+	66.123	1.423
σ	0.687	0.011	μ_v	13.089	0.263
κ_u	1.511	0.050	μ_u	154.224	8.638

Panel B: Summary Statistics

RMSE	1.974
Mean negative jump intensity	1.632
Mean negative jump size	-0.069
Mean positive jump size	0.015
Mean diffusive variance	0.054
Mean negative jump variance	0.015
Mean positive jump variance	0.006

Table 7: **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 statisticsfor 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	-1.000	0.012	c_0^+	9.666	0.471
$\overline{v}$	0.024	0.000	λ^-	18.967	0.189
κ	11.653	0.146	λ^+	68.833	1.148
σ	0.741	0.010	μ_v	32.642	0.690
κ_u	0.262	0.033	μ_u	44.814	6.474

Panel B: Summary Statistics

RMSE	2.050
Mean negative jump intensity	1.791
Mean negative jump size	-0.053
Mean positive jump size	0.015
Mean diffusive variance	0.055
Mean negative jump variance	0.010
Mean positive jump variance	0.004

Table 8: **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 statisticsfor the daily series of model-implied jump and variance estimates.Variances and jump intensity are given in annualized decimal units.

Panel A: Parameter Estimates

Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-0.752	0.010	c_0^+	4.397	0.445
$\overline{v}$	0.020	0.000	λ^-	24.316	0.425
κ	12.394	0.319	λ^+	66.762	2.024
σ	0.712	0.013	μ_v	27.147	0.579
κ_u	1.724	0.090	μ_u	415.332	41.169

Panel B: Summary Statistics

RMSE	1.835
Mean negative jump intensity	2.649
Mean negative jump size	-0.041
Mean positive jump size	0.015
Mean diffusive variance	0.032
Mean negative jump variance	0.009
Mean positive jump variance	0.002

Table 9: **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 statisticsfor 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	-1.000	0.014	c_0^+	5.099	0.979
$\overline{v}$	0.028	0.000	λ^-	12.688	0.389
κ	5.381	0.138	λ^+	66.693	5.117
σ	0.546	0.008	μ_v	30.845	1.347
κ_u	1.191	0.117	μ_u	25.114	15.159

Panel B: Summary Statistics

RMSE	2.164
Mean negative jump intensity	0.604
Mean negative jump size	-0.079
Mean positive jump size	0.015
Mean diffusive variance	0.040
Mean negative jump variance	0.008
Mean positive jump variance	0.002

Table 10: **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 statisticsfor 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	-0.973	0.038	c_0^+	1.564	0.250
$\overline{v}$	0.029	0.000	λ^-	10.823	0.284
κ	7.111	0.201	λ^+	24.902	1.458
σ	0.501	0.018	μ_v	24.804	1.026
κ_u	0.512	0.078	μ_u	20.066	6.613

Panel B: Summary Statistics

RMSE	2.867
Mean negative jump intensity	0.807
Mean negative jump size	-0.092
Mean positive jump size	0.040
Mean diffusive variance	0.063
Mean negative jump variance	0.014
Mean positive jump variance	0.005

Table 11: **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 statisticsfor 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	0.009	c_0^+	1.500	0.350
$\overline{v}$	0.040	0.001	λ^-	17.339	1.104
κ	6.530	0.239	λ^+	23.695	2.251
σ	0.683	0.012	μ_v	6.147	1.835
κ_u	1.244	0.438	μ_u	180.477	119.791

Panel B: Summary Statistics

RMSE	2.085
Mean negative jump intensity	3.419
Mean negative jump size	-0.058
Mean positive jump size	0.042
Mean diffusive variance	0.060
Mean negative jump variance	0.023
Mean positive jump variance	0.005

Table 12: **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 statisticsfor the daily series of model-implied jump and variance estimates.Variances and jump intensity are given in annualized decimal units.