The Pricing of Tail Risk: Evidence from International Option Markets*

Torben G. Andersen† Nicola Fusari‡ Viktor Todorov§

January 10, 2016

Abstract

The paper explores the global pricing of tail risk as manifest in equity-index options. We document the presence of a left tail factor that displays large persistent shifts which are not closely related to the corresponding dynamics of the return volatility factors. Moreover, the left tail factor is a potent predictor of future excess equity-index returns, while the volatility factors only forecast future equity variation, but not the expected returns. As such, the option surface dynamics embed separate risk and risk premium factors.

The evidence is robust across all indices explored, and the systematic deviations among the countries contain information about the differential risk and risk pricing during the financial and European sovereign debt crisis. We further explore the left tail risk premium dynamics across the indices. Again, the evidence is strikingly robust across the indices, but some economically important discrepancies appear, as the Southern European indices in our sample are subject to much higher risk pricing following the onset of the sovereign debt crisis in Europe. Interesting differences between Italy and Spain emerge, possibly speaking to the extent the investors in the two countries were differential affected by fears associated with the potential collapse of the euro currency area.

Preliminary and Incomplete

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

JEL classification: C51, C52, G12.

*Andersen gratefully acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The work partially supported by NSF Grant SES-1530748.

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

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

§Department of Finance, Kellogg School of Management, Northwestern University,1 Evanston, IL 60208; e-mail: v-todorov@northwestern.edu.

1 Introduction

The last decade has witnessed a great deal of turmoil in global equity markets. These events present a major challenge for modern asset pricing models. Does our standard framework capture the nature of risk pricing within individual equity markets? To what extent are the markets integrated?

In parallel, trading in the equity-index option markets across the developed world has grown sharply, with both more strikes per maturity and additional maturities on offer. In particular, the trading of options with short tenor has increased dramatically. As a result, we now have increasingly liquid and active trade prices for financial securities which reflect the pricing of financial market risk in each country. In this paper, we are able to draw on daily observations for option indices in the US (S&P 500), Euro-zone (STOXX), Germany (DAX), Switzerland (SMI), UK (FTSE), Italy (MIB), and Spain (IBEX) over 2007-2014.

The confluence of periods of turbulence render the last ten years an excellent "laboratory" for analysis of the way investors treat evolving financial risks and especially their attitude towards tail risk. Over our sample, three major shocks roiled the global financial markets, and we exploit the options data to study how tail risk was perceived across these episodes. At the same time, several of the European countries had unique exposures to the sovereign debt crises. It is of separate interest to determine whether these risks were priced similarly across the diverse equity indices. Finally, by combining the pricing of financial risks with ex post information on actual realized returns, return volatilities and jumps, we can gauge the size of the risk premiums and what factors drive the risk compensation. The striking heterogeneity in the performance over the sample, with the German market appreciating by an average of close to 5% per year and the Italian index depreciating by 10% annually, provides a wide range of experiences that may be associated with distinct dynamics of the respective option surfaces. In turn, these options should provide critical information about the expectations of the future distribution of the index returns as well as the pricing of the risk factors.

Standard option pricing models capture the dynamics of the equity-index option surfaces through the evolution of state variables that represent the volatility process for 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, are informative regarding the pricing of equity return risk, see, e.g., Andersen et al. (2015a). Hence, a distinct factor is necessary to account for the priced downside risk in the option surface relative to the regular volatility factors. In empirical work focused on the US markets, this state variable has emerged as a crucial equity risk premium

factor, see Andersen et al. (2015c).[1]

In comparing the standard two-factor option pricing models to a simplified two-factor version of the Andersen et al. (2015c) model, we find the latter to dominate uniformly in terms of fit to the option surfaces and to the spot volatility inferred from contemporaneous high-frequency equity-index data. This reflects a clear divergence between the time series evolution of priced tail risk and the level of market volatility for every option market we analyze. Thus, our analysis is built around a simple dynamic risk-neutral model with a separate volatility and left tail factor. A common feature emerges across the equity indices in the aftermath of crises: while the left tail factor is linked to volatility, it typically remains elevated long after market volatility subsides to pre-crisis levels. This is in line with the original Andersen et al. (2015c) evidence, but runs counter to the usual approach for modeling jump risk in the literature.

This stark separation of tail risk factor and market volatility has implications for the pricing and dynamics of market risks. We find the market volatility to be a strong predictor of future market risks, such as the jump intensity and the overall return variation. On the other hand, the component of the tail factor that is unspanned by the volatility factor, the "pure tail factor," predicts the future equity risk premium, while the volatility factor provides no auxiliary forecast power. The tail factor also is the primary driver of the negative tail risk premium, suggesting this is the operative channel through which it forecasts the equity risk premium. In particular, following crises, the asset prices are heavily discounted and the option-based tail factor is elevated long after market volatility resides. This combination turns out to be a strong signal that the stock market will rebound in the future.

In terms of return volatility, we find the European and US markets to evolve in near unison through the financial crisis of 2008-2009. In contrast, we observe signs of clear divergences in volatility during and after the initial European sovereign debt crisis. Overall, the UK, Swiss and, to some extent, the German volatilities remain very close throughout the sample. The largest divergences in the volatility dynamics occur between the US, UK and Swiss indices on one side and the Spanish and Italian ones on the other, with the latter representing the Southern European countries in our sample.

For the left tail factor, there are interesting commonalities and telling differences. Again the divergences are primarily associated with the Southern European indices. But whereas those two indices were qualitatively similar in their volatility responses, there is a stark discrepancy in the pricing of tail risk in Italy and Spain during the European debt crisis. Specifically, the tail factor of

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

the Spanish option market reacts strongly to either sovereign debt crisis, while the Italian response is much more muted, especially for the first of the two crises. Of course, the dramatic elevation in Spain may reflect expectations of much stronger future negative return shocks.

By exploiting the high-frequency equity-index return data, we also gauge the size of the left tail risk premium across the indices. We find a great deal of commonality in the tail pricing across the various indices, with Italy and Spain again standing out. In particular, we are able to largely dispel the hypothesis that the Spanish left tail factor is extraordinarily elevated relative to the Italian, simply because of the prevalence of large future negative return shocks in Spain. In other words, it seems that the Spanish market was gripped by "fears" and outsized risk premiums that were not present, even in Italy. One potential explanation is the possible rupture of the eurozone area, with a subsequent creation of a "Southern euro" that may include Spain, but likely not Italy. More formal analysis of such questions falls outside the scope of the current paper. It suffices to note that our analysis of the option surfaces brings the differential pricing of tail risk to the fore, setting the stage for more concrete analysis of such topics.

In summary, we find an overwhelming degree of robustness in the pricing of financial market risks across the US and the European indices. Moreover, the compensation for such risks appear to evolve similarly across the countries, with the exception of the two Southern European nations. The result confirm that the option surfaces allow for separate identification of risk and risk premium factors, represented by the spot variance and left tail factors respectively. For all indices, we verify the strong 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 separation between the tail and volatility factors is lacking and the associated risk premiums are intertwined, and thus not suitably identifiable.

The rest of the paper is organized as follows. We start our analysis with summarizing our key empirical results for the US market index in Section 2. Following that we review the data used in the paper in Section 3. Section 4 presents the model we use to fit the option surfaces and extract information from them. In Section 5 we review the estimation method. Section 6 contains our results for the risks and risk premia on international stock markets. Section 7 concludes. Additional details on various aspects related to the data, the estimation and the results are contained in Section 8.

2 Illustrative Evidence for the S&P 500 Index

This section exemplifies some of our key findings by reviewing results obtained from equity-index options at the Chicago Board Options Exchange (CBOE) written on the S&P 500 index as well as high-frequency returns on the e-mini S&P 500 futures at the CME Group. It extends evidence from Andersen et al. (2015c), but involves a simplified two-factor parametric model for the risk-neutral distribution. Moreover, the 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 existing evidence and serve as a benchmark for the European indices explored later on.

Our sample covers January 2007 – December 2014, matching the period for which we have a broad set of reliable option quotes available from the European countries. The risk-neutral return dynamics is governed by a traditional volatility factor and a jump factor with differential intensities and amplitudes for positive and negative jumps. Finally, the negative return jumps coincide with positive volatility jumps. Even if the model is significantly pared down relative to Andersen et al. (2015c), it captures the salient features of the option surface and equity-index return dynamics across time and for all the individual markets. In particular, the model fits the option prices significantly better than the traditional two-factor stochastic volatility jump-diffusive specification. Additional details regarding the data sources, the specific parametric models used, and the estimation and inference techniques applied are deferred to subsequent sections.

2.1 The U.S. Evidence

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

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

Using our parametric model for the risk-neutral dynamics, we extract the implied realization of

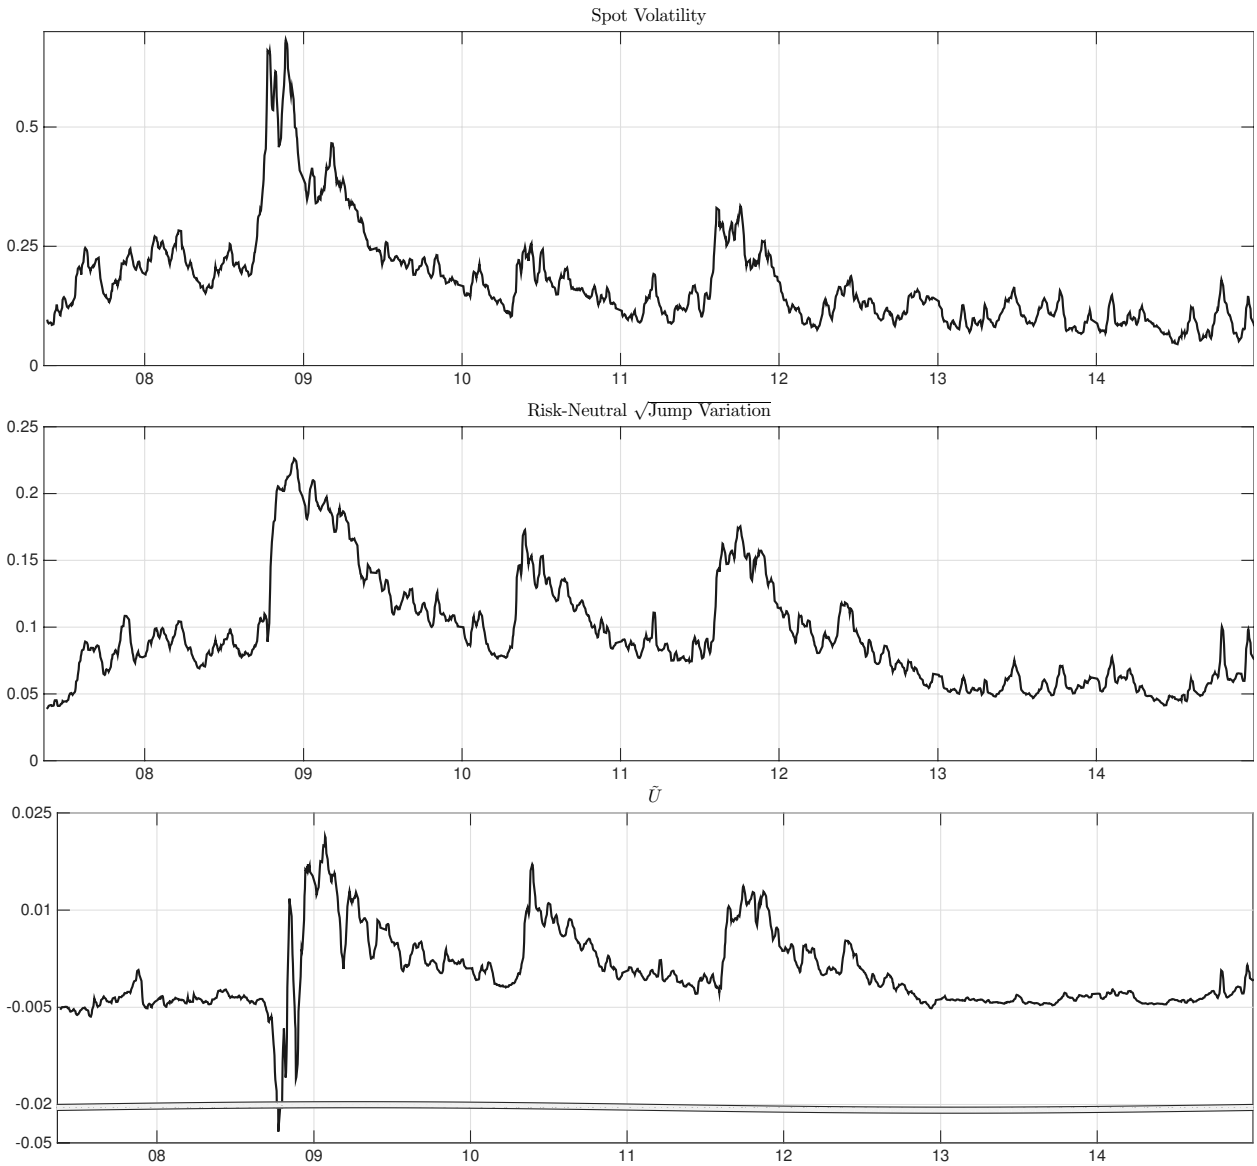

Figure 1: **Implied Spot and Negative Jump Volatilities**. The figure displays the five-day trailing moving average of the spot volatility (top panel), the (square-root of the) negative jump variation (middle panel), and the residual of the negative jump variation regressed on the spot variance (bottom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

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

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

In the bottom panel of Figure 1, we display the residual from the regression of the negative jump variation on the spot variance which we denote with $\widetilde{U}$. 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 heterogeneous 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.

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

In our reduced form model, two state variables – the spot variance and the 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 distribution for the actual equity-index returns and, in particular, the expected excess returns and risk. In affine models, this relation is linear. We first consider the risk premium relation. Hence, Figure 2 depicts the significance of

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

the regression coefficients and the degree of explanatory power of the implied spot variance, V_t, and jump intensity, U_t, for the future excess returns on the S&P 500 index through a simple bivariate regression. Given the relatively short sample size, we limit ourselves to the horizons spanning one week to about half a year.

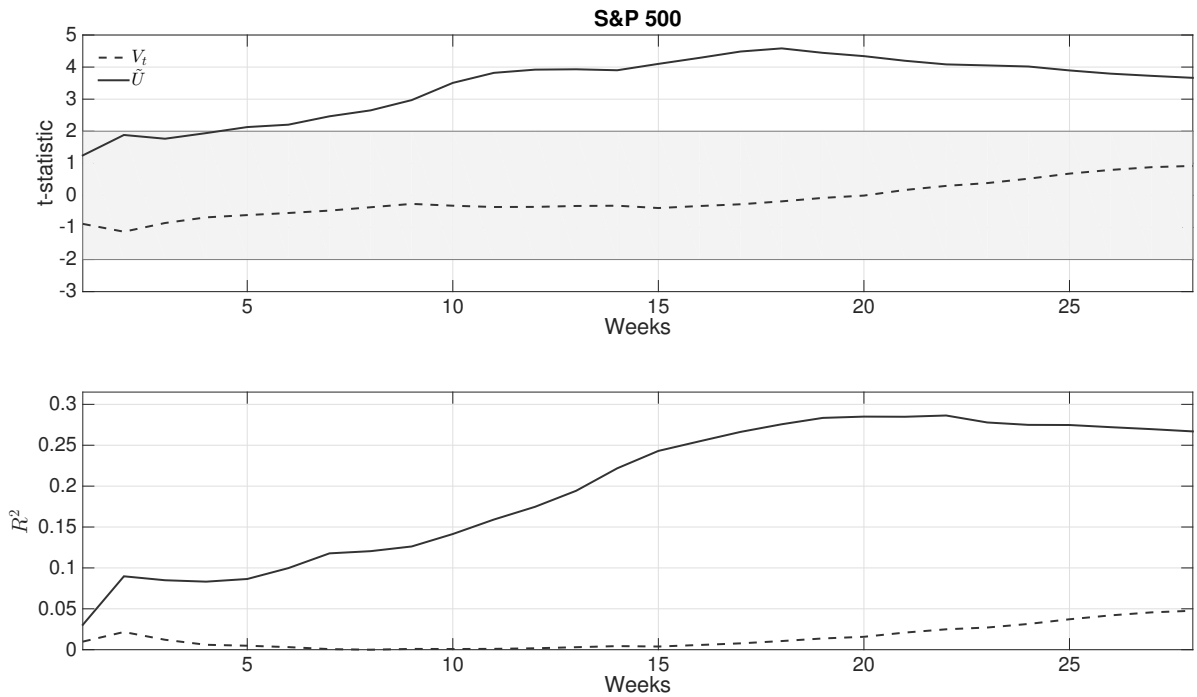

Figure 2: **Weekly Predictive regression on excess returns.** The figure reports findings stemming from a regression of the future cumulative excess weekly equity-index returns on the current option-implied state variables, i.e., the spot variance and the negative jump intensity. The top panels displays the t-statistic for the regression coefficients on the state variables, while the bottom panel depicts the corresponding regression R^2.

We find the excess returns to be highly significantly linked to the component of the left jump tail intensity factor unspanned by V_t, i.e., $\widetilde{U}_t$ for horizons beyond two months, with a maximum attained around four months. The associated degree of explained variation in the excess returns exceeds 25% for horizons beyond 3 1/2 months. At the same time, we note that the implied spot volatility is insignificant, so the information regarding the pricing of equity return risk is fully contained in the negative jump factor.

In contrast, when we regress the future (realized) return variation, stemming from both diffusive volatility and jumps, or just the future diffusive return variation, on the state variables, the relative explanatory power is entirely reversed. The jump intensity unspanned by the market volatility has

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

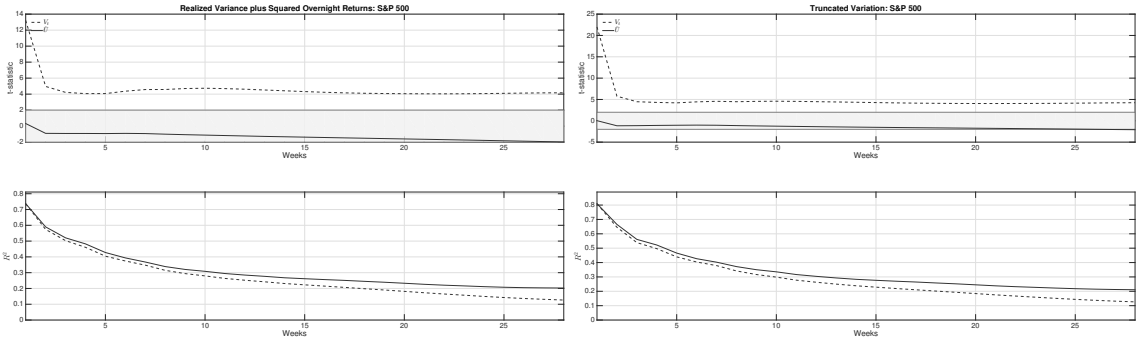

Figure 3: **Predictive regression on return variation measures**. The figure reports findings stemming from regressions of future cumulative return variation measures on current option-implied state variables, i.e., the spot variance and the negative jump intensity. The top panels displays the t-statistic for the regression coefficients on the state variables, while the bottom panel depicts the corresponding regression R^2. The left panels concern the total realized variation, accounting for the diffusive variation, the squared jumps and the squared overnight (close-to-open) returns, whereas the right panels refer to the diffusive return variation alone, estimated by the Truncated Realized Variation, TRV, estimator.

These findings raise a number of questions. Are similar option pricing and risk premium patterns present in other developed economies? Are the fundamental U.S. state variables related to corresponding factors in other countries? That is, 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 provide a more formal review of our risk-neutral two-factor model.

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. The database also provides detailed information regarding the features of the various equity-index option contracts as well as the zero curve for the relevant currency. In order 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 international indices: USA (SPX), Europe (SX5E), Germany (DAX), Switzerland (SMI), United Kingdom (UKX), Italy (MIB), Spain (IBEX). For each index, in Table 1, we provide the exchange trading hours, which are used 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. It reports either the end-of-the-day quotes or the exchange settlement price for each contract. It is not possible to distinguish the two, but the vendor reports that 98% of the data represents settlement prices and only 2% reflect trade prices, with some variability depending of the specific exchange.

We create the final option sample through the following steps. First, for each option maturity we compute the associated 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 cross sections with at least 5 put-call contracts with the same strike price, and then extract the futures price exploiting the full set of option pairs with the same strike. Third, we apply a few filters to ensure that the prices are reliable: we only use options with a tenor below one year, as longer maturity contracts tend to be illiquid. However, contrary to the prior literature, we include very short-maturity options in our analysis. This is due to the recent successful introduction of short-dated options by several exchange worldwide. These options are particularly informative regarding the current state of the return dynamics, see, e.g., Andersen et al. (2015b) 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 prices 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 quadratic variation (QV), the truncated variation (TV), and the jump variation (JV) following the procedure of Bollerslev and Todorov (2014) and Andersen et al. (2015c), see also the appendix for further details. Moreover, we construct a measure of the *local* continuous variation (LV) – relying on the returns over a three hour window prior to the close of the option market – which is used in our estimation procedure. Table 2 reports, for each index, the country, the associated exchange, and various contractual details. 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 trading day. [3]

Table 3 provides summary statistics for the daily measures, obtained from the high-frequency index returns. The vast divergence in the fortunes of these indices is striking. While the U.S. market experiences annual returns in excess of 4.5%, and the German index does even 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" where volatility rises following negative return shocks – 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 least exposed to this type of negative shocks, while the differences across the European indices are minor. However, the relative jump counts are somewhat skewed by the fact that the U.S. markets are closed during the early trading in Europe, when many jumps may have materialized, but simply could not be observed in the U.S. The split into positive and negative return jumps reveal that the indices associated with the euro-zone had substantially more downward than upward jumps, possibly reflecting the impact of the sovereign debt crises. In contrast, the jump direction is nearly symmetric for the remaining indices. The stale price statistic reveals that the Swiss, and possibly also the euro-zone STOXX, index may be impacted by illiquidity, inducing a potential downward

[3] The trading hours are slightly ambiguous, as electronic trading may occur outside the stated interval in the table. For example, the S&P 500 e-mini futures trade almost 24 hours on the GLOBEX platform, while the indicated period refers the time when pit trading is active along with the electronic trading. The table follows the conventions adopted by TickData.

bias in the realized variation measures, while the remaining indices are highly liquid. However, the closeness of the return standard deviation measure obtained from the high-frequency and daily returns suggest that the general impact of illiquidity on the measures is trivial.

Hence, the realized equity risk premiums deviate greatly across the indices in our sample. At the same time, the indices respond quite 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 the Appendix.[4] 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. Of course, these quantities jointly impact the individual equity and variance risk premiums over the sample whose realizations are hugely heterogeneous. 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 allows the shape of the option surface to speak more cleanly to the concurrent pricing of equity risk.

[4]The daily return correlations must be interpreted with some care, as they are not fully synchronized, especially for the U.S. and the European indices. Thus, the numbers in the table should be viewed as lower bounds on the true (synchronized) correlations.

Index Name	Country	Exchange	Ticker	OptionMetrics	Date Starts	Trading Hours (Exchange)	Tick Size	Multiplier
				North America				
S&P 500	USA	CBOE	SPX	SPX	1996	8:30 a.m. - 3:15 p.m.	0.05 Points	100 $
				Europe				
DJ EURO STOXX 50	Europe	EUREX	OESX	SX5E	2002	8:50 a.m. - 5:30 p.m.	0.1 Points	10 €
DAX	Germany	EUREX	ODAX	DAX	2002	8:50 a.m. - 5:30 p.m.	0.1 Points	5 €
SMI	Switzerland	EUREX	OSMI	SMI	2002	8:50 a.m. - 5:20 p.m.	0.1 Points	10 CHF
FTSE 100	UK	EURONEXT	ESX	UKX	2002	8:00 a.m. - 4:30 p.m.	0.5 Points	10 £
FTSE MIB	Italy	IDEM	MIBO	MIB	2006	9:00 a.m. - 5:40 p.m.	1.0 Points	1 €
IBEX 35	Spain	MEFF	IBX	IBEX	2006	9:00 a.m. - 5:35 p.m.	1.0 Points	2.5 €

Table 1: **Option contract specifications**. For each option contract we report the underlying index, the corresponding country, the name of the exchange, the ticker symbol, the symbol in the OptionMetrics database, the earliest date from which data are available, and finally the trading hours, the tick size and the multiplier (as of December 2014).

Index Name	Country	Exchange	Ticker	TickData	Data Starts	Daily Trading Hours (Exchange)	Tick Size	Multiplier
				North America				
S&P 500 E-mini	USA	CME	ES	ES	1997:9	8:30 a.m. - 3:15 p.m.	0.25 Points	50 $
				Europe				
DJ EURO STOXX 50	Europe	EUREX	FESX	XX	1998:7	8:00 a.m. -10:00 p.m.	1.0 Points	10 €
DAX	Germany	EUREX	FDX	DA	1997:1	8:00 a.m. -10:00 p.m.	0.5 Points	25 €
SMI	Switzerland	EUREX	FSMI	SW	1998:10	8:00 a.m. -10:00 p.m.	1.0 Points	10 CHF
FTSE 100	UK	EURONEXT	FFI	FT	1998:7	8:00 a.m. - 9:00 p.m.	0.5 Points	10 £
FTSE MIB	Italy	IDEM	SPMIB	II	2004:10	9:00 a.m. - 5:40 p.m.	0.5 Points	10 £
IBEX 35	Spain	MEFF	MFXI	IB	2003:7	9:00 a.m. - 8:00 p.m.	0.5 Points	10 £

Table 2: **Futures contract specifications**. For each index futures contract we report the country, the option exchange, the ticker symbol, the TickData symbol for the contract, the earliest date for which data are available, the daily trading hours, the tick size, and the multiplier (as of December 2014).

	ES	XX	DA	SW	FT	II	IB
$\sqrt{RV}$	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
LJV/JV	0.49	0.58	0.56	0.56	0.52	0.56	0.52
N. Neg. Jumps	5.60	28.81	29.49	9.01	14.17	25.46	27.04
N. Pos. Jumps	6.61	21.79	21.49	7.74	12.02	17.69	17.53
Average Daily Return (%)	4.58	-3.57	4.81	0.19	0.47	-9.99	-4.07
Std 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 3: **Summary statistics for equity-index futures.** The numbers are annualized and given in percentage form, except for the ratios in rows 3-4, the number of jumps in rows 5-6, and the last row dealing with stale prices. The realized variation measures (RV and TV) are computed within the trading day using log returns, then average and scaled to represent one calendar year. We report the square-roots of these measures to cast them in annualized standard deviation units. The number of jumps is annualized and concerns only jumps within the trading days. The realized variation measures, RV and TV, and the number of jumps are obtained via the procedure detailed in the appendix. The daily standard deviation (std) is computed from the index close-to-close return. The stale prices are obtained as the percentage of 5-minute intervals for which neither the bid nor ask price changed across the five observations obtained by retaining the final quotes in each calendar minute within a given 5-minute period.

4 Model

We denote a generic equity market index price with X. Our two-factor model for the risk-neutral equity index dynamics is given by the following restricted version of the representation in Andersen et al. (2015c),

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

where $(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}})$ is a two-dimensional Brownian motion with $\text{corr}\left(W_t^{\mathbb{Q}}, B_t^{\mathbb{Q}}\right) = \rho$. In addition, μ is 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 measure, 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 volatility and (negative) jump intensity are co-linear, albeit with distinct proportionality factors, μ_v and μ_u. This specification involves a substantial amplification from the negative price shocks to the risk factors. The compensator characterizes the conditional jump distribution and is given by,

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

The right hand side refers to negative and positive price jumps, respectively. Following Kou (2002), we assume that the price jumps are exponential, with separate tail decay parameters, λ_- and λ_+, for negative and positive jumps. Finally, the jump intensities are governed by the $c^-(t)$ and $c^+(t)$ coefficients which evolve as affine functions of the state vector,

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

This representation involves a large set of parameters. At the estimation stage, we zero out insignificant coefficients. Specifically, for the negative jump intensity, we fix c_0^- and c_v^- to zero and, for identification, normalize c_u^- to unity, so the value of U is identical to 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 GARCH models. Third, the jump intensity is decoupled from the volatility. This is unlike most existing option pricing models in the literature, with the notable exception of Christoffersen et al. (2012) and Li and Zinna (2015). Nonetheless, as stressed by Andersen et al. (2015a), the model (1) still belongs to the affine class of models of Duffie et al. (2000).

5 Estimation Procedure

We follow the estimation and inference procedures developed in Andersen et al. (2015a). The option prices are converted into the corresponding Black-Scholes implied volatilities (BSIV), i.e., any out-of-the-money (OTM) option price observed at time t with tenor τ (measured in years) and log moneyness $k = \log(K/F_{t,t+\tau})$ is represented by the BSIV, $\kappa_{t,k,\tau}$. For a given state vector, $\mathbf{S}_t = (V_t, U_t)$, and risk-neutral parameter vector θ, the corresponding model-implied option price is given by $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$. Estimation of the parameter vector and the period-by-period realization of the state vector now proceeds by minimizing the distance between the observed and model-implied BSIV across the full sample in a metric that also penalizes for the discrepancies between the inferred spot volatilities and those estimated from high-frequency return observations on the underlying asset, $\sqrt{\widehat{V}_t^n}$. The latter are obtained from five-minute returns over a three-hour window prior to the close of the trading day using the so-called truncated realized volatility estimator, as implemented in Andersen et al. (2015c). 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 κ_{t,k_j,τ_j} denoting the BSIV for option j observed at date t, log-moneyness k_j and tenor t_j. Hence, $k_j \leq 0$ indicates OTM put options and $k_j > 0$ OTM call options. These index options are more liquid and have lower bid-ask spreads than in-the-money options. Put-call parity implies that we still cover the full range of strikes relevant for extracting information about the pricing of future contingencies related to the state of the equity market. For each option in our sample, the

corresponding model-implied BSIV is denoted $\kappa_{k_j,\tau_j}(\mathbf{S}_t,\theta)$.

We obtain point estimates for the parameter vector θ and the period-by-period state vector $\mathbf{S}_t = (V_t\,, U_t\,)$ from the following optimization problem,

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

where the penalty for the deviation between the realized and model-implied spot volatility is given by $\xi_n > 0$ and the superscript n denotes the number of high-frequency returns exploited by the spot return variance estimator, $\widehat{V}_t^n$. For our implementation with a given fixed n, we set $\xi_n = 0.05$, as in Andersen et al. (2015c). Moreover, to reduce the computational burden, we only estimate the system for options sampled on Wednesday or, if this date is missing, the following trading day. The critical feature ensuring good identification of the parameters is to obtain observations across heterogeneous constellations of the option surface. We achieve this by sampling widely across the full sample period. The shape of the surface varies dramatically across the early and late years, when the market is fairly quiet, relative to the periods associated with the financial and European debt crises. Once the parameter vector and the state variable realizations for those Wednesdays have been obtained, it is straightforward to "filter" the state variables for the remaining trading days, exploiting the estimated parameters and the criterion (4). Thus, we have daily estimates for the state realizations available, even if full-fledged estimation is performed only for weekly data,

6 Estimation Results

Given our limited sample period and the somewhat lower number of observations available for some of the European indices relative to the S&P 500 illustration in Section 2, our purpose is not to obtain a perfect option pricing model, but rather to settle on a specification that captures the salient features across all indices in a robust manner. Thus, our specification involves only a single volatility factor and the dynamics of the second, negative jump intensity, factor is pared down and simplified relative to the modeling in Andersen et al. (2015c). This largely eliminates instances where separate identification of the two factors is troublesome and ensures that the comparisons across the countries are meaningful. We start this section with documenting the extent to which the current model improves on standard option pricing models in which the factors driving the option surface are all volatility factors.

6.1 Comparison with the Standard Two-Factor Stochastic Volatility Model

The standard two-factor stochastic volatility option pricing model, used extensively in prior work, see e.g., Bates (2000, 2003) , is one in which the diffusive volatility is a sum of two factors V_1 and V_2 which are jointly affine, and in which the jump intensity is an affine function of these two volatility factors. To conserve space, the formal model is relegated to the Appendix. Here, we simply note that the two alternative models for the pricing options have the identical number of parameters and both feature two state variables to capture the dynamics of the option surfaces. Therefore, any differences in the models' performance may be attributed solely to their relative ability to fit the option surfaces.

The key distinguishing feature of the standard two-factor volatility model (2FV), from the one we are using here (2FU) to unlock the information in the option surfaces, is the tight link it implies between the level of market volatility and the (risk-neutral) jump intensity. This connections is loosened markedly in our model which, on the other hand, has less flexibility in capturing the presence of distinct volatility components. We compare the models' ability to fit the option surface and to match nonparametric high-frequency based volatility estimates, i.e., the values of the two components of the objective function in (4) at their parameter and state vector estimates. If the volatility surfaces convey separate critical information regarding the pricing of jump risk and this is more important than a refined modeling of the volatility dynamics, then we would expect our 2FU model to dominate the traditional (2FV) specification. Alternatively, if refined modeling of volatility is key to accommodate the risk dynamics and pricing of risk, then the traditional (2FV) models should prevail. The results are reported in Table 4. Starting with the S&P 500 index, we find that the overall reduction in the value of the objective function in favor of our (2FU) model is a nontrivial 30%. The improvement stems both from the fit to the option surface and in matching the model-free estimate of spot volatility. It turns out that this conclusion is universal for all the equity index options in our sample. The biggest improvement in model performance is for the Euro STOXX index where the reduction in the objective function is an impressive 43%. The smallest improvement is for the Spanish market index of 9%. This suggests that the option data are highly informative regarding the connections between volatility and jump risks. For all the option markets it is highly desirable to to detach the risk-neutral jump intensity from the level of market volatility. Indeed, as for the S&P 500 index, we demonstrate below that this separation of jump intensity from volatility enables the two factors to attain critically distinct roles in the characterization of the risk and risk premium dynamics.

	RMSE(IV)		RMSE(Vol)		Obj. Function	
	2FV	2FU/2FV	2FV	2FU/2FV	2FV	2FU/2FV
SPX	2.51	0.85	6.92	0.93	9.73	0.70
SX5E	2.60	0.73	5.77	0.90	9.64	0.57
DAX	2.44	0.76	5.35	0.89	8.46	0.62
SMI	1.72	0.94	3.34	0.89	4.56	0.88
UKX	2.35	0.86	4.40	0.89	7.49	0.72
MIB	3.53	0.79	5.68	0.75	16.73	0.62
IBEX	2.13	0.99	5.34	0.88	7.42	0.91

Table 4: **Models Comparison**. For each Index: the second column reports the RMSE (in implied volatility) in percentage for the 2FV specification, the third column reports the ratio between the RMSE (in implied volatility) of the 2FU and the 2FV. Similarly columns fourth and fifth show the RMSE (and their ratio) in terms of Variance fit, and finally columns sixth and seventh show the value of the objective function (and their ratio).

6.2 Country-by-Country Factor Realizations

We continue by directly comparing the option-implied spot variance factor for each of the individual equity indices enumerated in Section 3 to the spot variance extracted from the S&P 500 options, depicted in Section 2. Figure 4 provides the plot for Euro STOXX (Eurozone), DAX (Germany), SMI (Switzerland), FTSE 100 (U.K.), FTSE MIB (Italy), and IBEX 35 (Spain). The most striking feature is the extraordinary close association between some of these variance factors and the S&P 500 spot variance, not just in terms of correlation, but also level. For example, the U.K. variance factor is barely distinguishable from the S&P 500 factor throughout the sample, while the Swiss factor only deviates during a couple of episodes under the fixed Swiss franc-euro exchange rate policy implemented in September 2011 and lasting through the remainder of our sample. In the former case, the correlation between the variance factors is about 98% and in the latter around 95%. In contrast, the discrepancies between the S&P 500 and German DAX indices emerge during the second phase of European sovereign debt crisis, when the DAX volatility spikes significantly more than the S&P 500, and the (positive) gap remains for the remainder of the sample, albeit to a varying degree. For the broader euro-zone STOXX 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 and up to early 2010, there is a marked differential impact of the sovereign debt crisis. Moreover, the relative effect across the national indices appears consistent with the perceived sensitivity of the respective economies to the European crisis. This becomes even more transparent for the Italian and Spanish indices, as both react much more strongly to the crisis events, but with different amplitudes across the main episodes. For the latter indices, the volatility levels attain a plateau well above the others ever since the first clear signs of the sovereign debt problems surfaced in early 2010. This systematic divergence over the second half of our sample lowers the volatility correlations for MIB and IBEX with S&P 500 to 0.75 and 0.77, respectively.

Next, Figure 5 depicts the option-implied negative jump variation for these indices along with the corresponding quantity for the S&P 500. These jump variation series are directly proportional to the negative jump intensity factor, U, so they also provide a genuine comparison of the option-implied state variables across indices.[5]

Qualitatively, the pattern is similar to the one observed for the volatility factors as, for each index, the volatility and jump factors are highly correlated. Nonetheless, as for the U.S., the relative size of the spikes in the jump intensity versus volatility varies substantially, with the European crisis inducing a stronger surge in the jump intensity relative to (diffusive) volatility. Again, the U.K. and U.S. series evolve in unison and display a correlation of over 98%, even if the British jump variation is slightly lower throughout. Similarly, the jump intensity factor for the Swiss index correlates strongly with the S&P 500 factor, although the Swiss factor tends to be above the one in the U.S. before and during the financial crisis and then below after the summer of 2009. We also observe strong coherence between the S&P 500 series and the STOXX 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 STOXX 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 distinct phases of the debt crisis relative to the other countries or 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. Our decomposition of the primary risk

[5]The negative (risk-neutral) jump variation is formally defined as $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx)$, and in our model it is equal to $\frac{2}{\lambda_-^2} U_t$.

factors documents a substantially larger increase in return volatility for these two indices along with a further amplification of the negative jump risk.

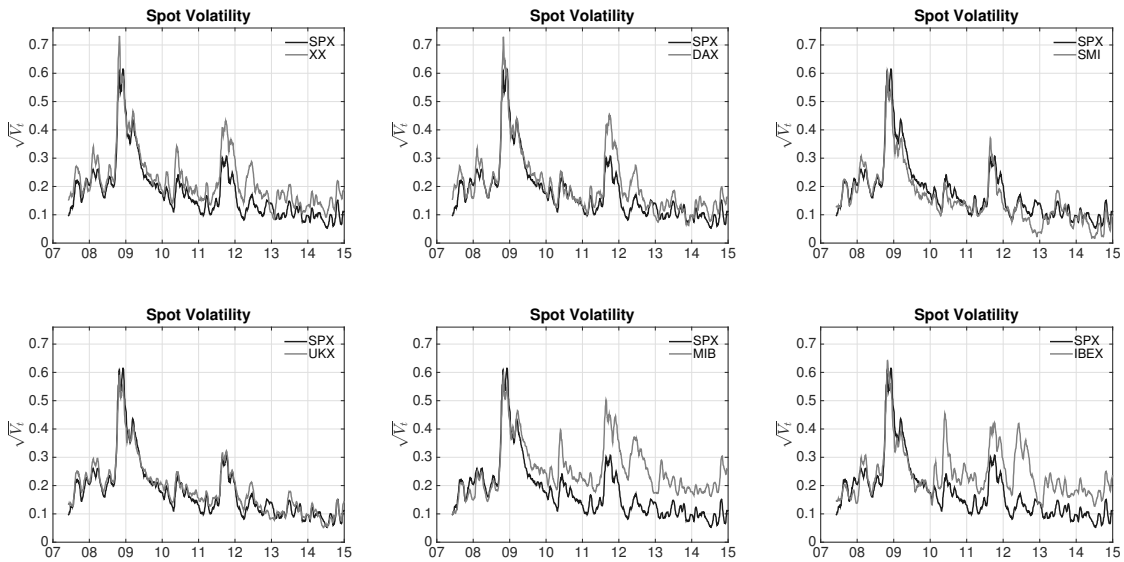

Figure 4: **Variance factor comparison.** For each option-implied spot variance factor, we report the trailing five-day moving average. The pairwise correlation between the variance factor for the S&P 500 and each of the European indices is as follows: SPX-XX: 0.948; SPX-DAX: 0.939; SPX-SW: 0.953; SPX-FT: 0.979; SPX-MIB: 0.750; and SPX-IBEX: 0.770.

The coherence across the volatility and jump variation series as well as the striking discrepancies observed during crisis episodes, lending themselves to ready economic interpretation, adds further credence to the robustness of our methodology in extracting the salient pricing features from the option surfaces.

6.3 Option-Based Prediction of the Equity Risk Premium and Return Variation

This section explores the ability of the option-implied factors, spot variance and left jump intensity, to forecast the (realized) equity returns and the (realized) future return variation (the sum of the squared returns). The former signifies whether the factors are associated with the equity risk premium, while the latter speaks to their predictive power for future equity risk, as captured by the magnitude of the ensuing return variation.

Specifically, we regress the future excess returns for each of the European equity indices on their option-implied state variables. We do this first for the local state variables in isolation, but then also in a regression augmented by the corresponding U.S. state variables. The latter

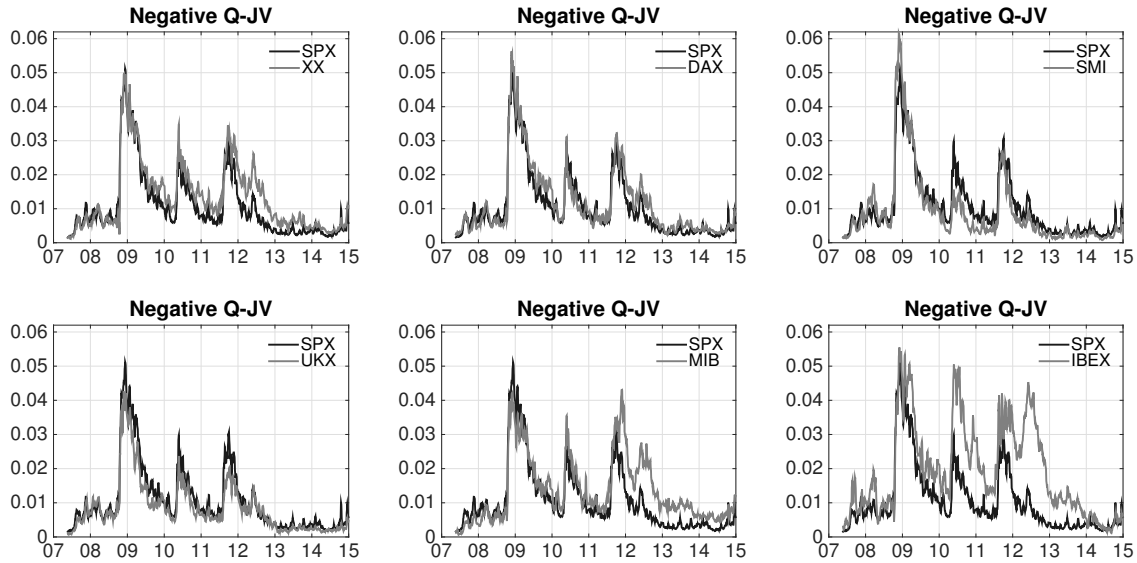

Figure 5: **U factors comparison**. For each option-implied negative jump intensity factor, we report the trailing five-day moving average. The pairwise correlation between the jump factor for the S&P 500 and each of the European indices is as follows: SPX-XX: 0.949; SPX-DAX: 0.970; SPX-SW: 0.951; SPX-UKX: 0.981; SPX-MIB: 0.854; and SPX-IBEX: 0.798.

specifications test for the impact of the state of the U.S. market, viewed as a potential proxy for international risk or pricing factors. In other words, is there evidence of integration across the equity-index option markets, with a common factor exerting a non-trivial impact on the risk-neutral dynamics in Europe? For ease of interpretation, we eliminate a large degree of collinearity among the explanatory variables through sequential orthogonalization. The single country regressions exploit the local spot variance factor and adds the component of the local left jump factor that is orthogonal to the spot variance as the second explanatory variable. When additional variables are used, all are included in a corresponding fashion, i.e., we only retain the variation in the subsequent factors that are orthogonal to the preceding explanatory variables in the regression. Hence, we allow the first factor to capture as much of the variation in the excess returns as possible, and then explore whether the additional factors provide significant incremental explanatory power.

Our first set of regressions take the form, for $t = 1, \ldots, T - h$,

$$r_{t,t+h} \;=\; p_{t+h} - p_t \;=\; c_0 \,+\, c_v \cdot V_t \,+\, c_u \cdot \widetilde{U}_t \,+\, \epsilon_t \,, \tag{5}$$

where $p_t = \log X_t$ represents the end-of-trading logarithmic index level, h is the horizon (in days), $r_{t,t+h}$ denotes the future h-day continuously-compounded return, and $\tilde{U}_t$ is the local option-implied left jump intensity orthogonalized with respect to the spot variance.

The second set of regressions explores whether the U.S. option-implied state variables may serve

as global factors. Thus, the initial regressor is now the spot variance extracted from the S&P 500 index options, followed by suitably orthogonalized S&P 500 jump intensity, and then the local spot variance and jump intensity factors.

$$r_{t,t+h} \;=\; c_0 \,+\, c_v^{us} \cdot V_t^{us} \,+\, c_u^{us} \cdot \widetilde{U}_t^{us} \,+\, c_v \cdot \widetilde{V}_t \,+\, c_u \cdot \widetilde{U}_t \,+\, \epsilon_t \,. \tag{6}$$

Given our limited sample period, we run the predictive regression on a weekly basis, forecasting from 1 to 28 weeks, or roughly 6 1/2 months, into the future. We compute robust Newey-West standard errors using twice the number of lags relative to forecasting horizon. For example, if we forecast the future 28-week excess return, we include 56 lags in the computation of the standard errors. 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 exclude days from the regression where one of the explanatory variables takes an absolute size beyond fifteen standard deviations of the median value. Importantly, we do not exclude the corresponding daily return or return variation from the cumulative measures appearing as left-hand-side variables, i.e., the extreme return or volatility realizations are included in the definition of the multi-horizon excess returns and return variation measures. The deletion of extreme regressors excludes less than 1% of the observations on average.

The regression results are, qualitatively, strikingly similar across countries. To conserve space, we report findings for four illustrative indices only and defer the remaining results to an appendix. We start with the EuroSTOXX index. It represents a scenario in-between Germany (DAX) on the one hand and Italy (MIB) and Spain (IBEX) on the other, although it turns out to generate results more closely aligned with the DAX rather than the latter two.

The top left panel of Figure 6 plots the t-statistics for the EuroSTOXX regression coefficients in equation (5), while the left bottom panel displays the corresponding regression R^2. Consistent with the hypothesis of a slowly moving equity risk premium, the predictive power of the regression rises steadily over time. From the latter panel, we see that the unpredictable return component dominates for short weekly horizons, but since random noise innovations cancel over time, the predictable component emerges clearly for the longer horizons. At about three months, the R^2 surpasses 10% and exceeds 15% after 6 months. From the top panel, it is evident that the explanatory power stems from the jump intensity factor and not the (diffusive) volatility. In fact, the second lower curve in the left bottom panel reveals that the volatility alone is largely immaterial and only provides marginal explanatory power after about half a year. Thus, the commonly employed volatility factor has no discernable relationship with the equity risk premium, while unrelated vari-

ation in the left side of the option surface is indicative of systematic shifts in the pricing of equity risk. Of course, this is also entirely consistent with our findings for the S&P 500 index in Section 2 as well as Andersen et al. (2015c).

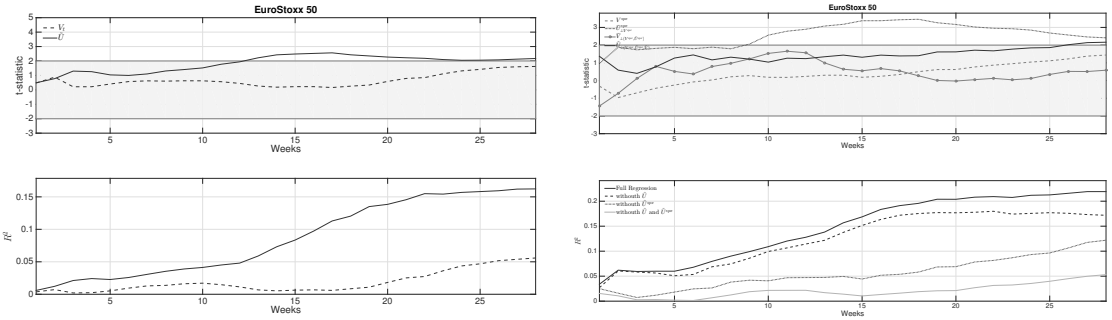

Figure 6: **Predictive Regressions, EuroSTOXX Excess Returns**. Left Panel: Index specific factors. Right Panel: Index specific plus US factors.

Turning to the right panels in Figure 6, we find the U.S. jump intensity factor playing a strikingly similar role for the Euro-zone excess equity returns as the corresponding local factor. In fact, there is no indication that the U.S. factors are less informative than the euro-zone implied factors themselves. In addition, both the U.S. and European spot variance factors are bereft of explanatory power for the future returns. In this sense, the two option markets appear highly integrated and the results are consistent with the presence of a global equity risk premium factor.

But what do the Euro-zone option-implied factors represent in terms of the risk characteristics of the underlying equity-index return dynamics? To address this issue, we again consider regressions of the form (5) and (6), but the dependent variable is now a measure of the future realized return variation, $RVM(t, t+h)$. The latter is constructed from high-frequency intraday observations on the equity-index futures, and possibly 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 equity index, see, e.g., Andersen et al. (2003).

$$RVM_{t,t+h} \;=\; k_0 \,+\, k_v \cdot V_t \,+\, k_u \cdot \widetilde{U}_t \,+\, \epsilon_t \,, \tag{7}$$

Ignoring the overnight (squared) returns in the construction of RVM implies better precision, but leaves out a non-trivial component of the return variation. Nonetheless, using either measure, the left panels of Figure 7 reveal, in line with the S&P 500 evidence in Section 2, that the left jump intensity factor has no explanatory power for the ex-post realized return variation. Instead, all predictor power is concentrated in the implied variance which is known to be a powerful predictor of short-term return volatility. The striking feature is the absence of any predictive power in the

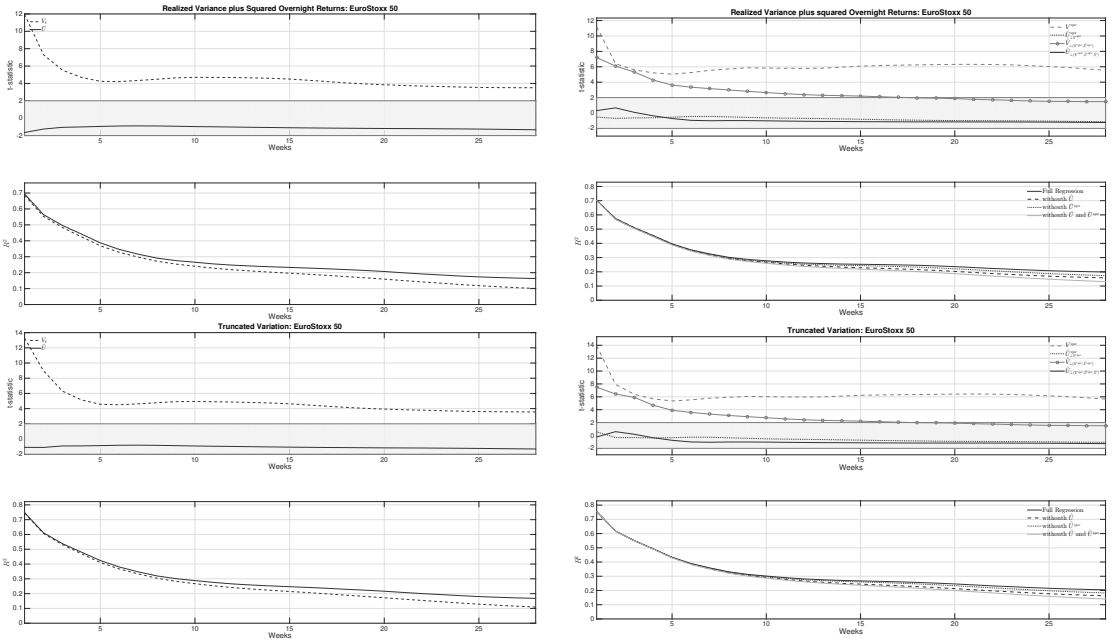

Figure 7: **Predictive Regressions for EuroSTOXX High-Frequency Return Variation Measures**. Top Panels: RV plus squared overnight returns. Bottom Panels: Truncated return variation. Left Panels: Index specific factors. Right Panels: Index specific plus US factors.

jump factor. After all, this is the component that is relevant for forecasting the equity returns. As in Section 2, we are led to the conclusion that our two state variables, disentangled from the option surface dynamics through our parametric representation for the risk-neutral return dynamics, generate a clean separation between risk (V) and risk premium (U) factors.

In the right panels of Figure 7, we focus on the predictive power of the U.S. state variables for the euro-zone equity return variation. Again, we find the jump intensity factors to be irrelevant, while the spot variance factors are highly significant. However, overall, the explained variation is very similar to before, implying that the U.S. factor is highly correlated with the corresponding euro-zone factor but carries little incremental information about the local return variation.

We now review the identical regression evidence for the Swiss index, SMI. For the excess return regression, we observe one noteworthy difference. The explanatory power of the U.S. jump intensity factor is considerable higher for the future swiss excess returns than the local Swiss jump intensity factor. As before, the volatility factors have no meaningful predictive power.

For the future return variation of the Swiss index, the jump variation factors are insignificant, while the two spot variance factors provide a high degree of forecast power. As before, the incremental forecast power of the U.S. variance factor is limited.

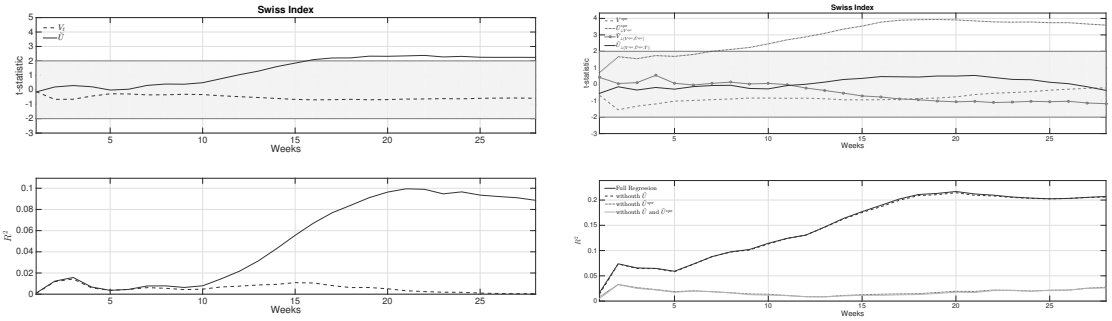

Figure 8: **Predictive Regressions, SMI Excess Returns**. Left Panel: Country specific factors. Right Panel: Country plus US factors.

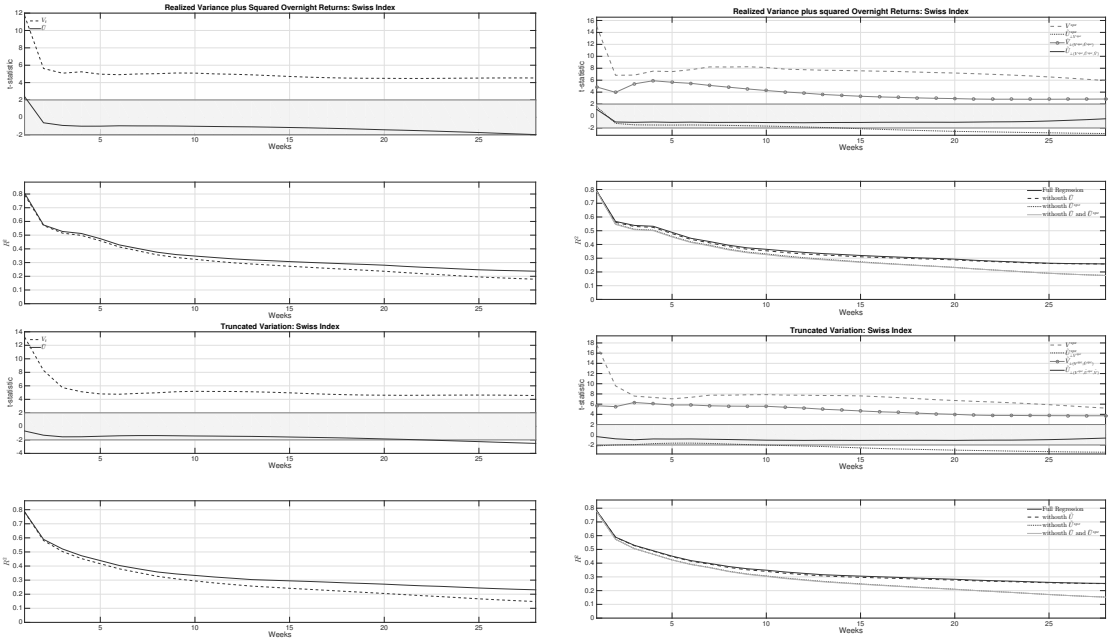

Figure 9: **Predictive Regression, SMI Return Variation Measures**. Top Panels: RV plus squared overnight returns. Bottom Panels: Truncated Variation. Left Panels: Country specific factors. Right Panels: Country and US factors.

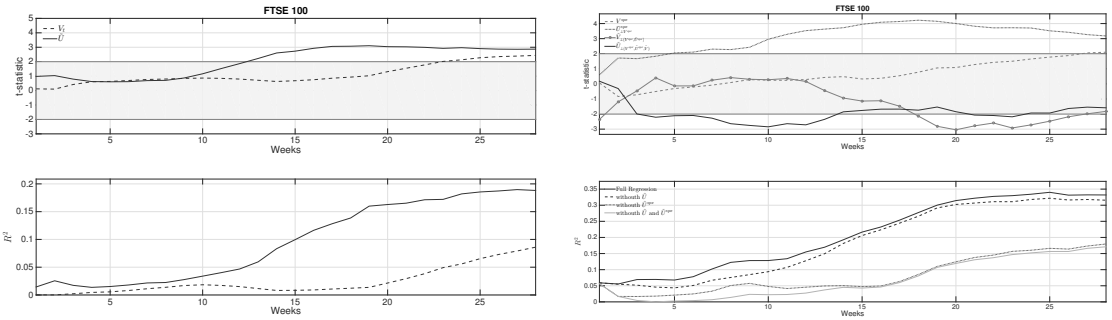

Figure 10: **Predictive Regression, FTSE Excess Returns**. Left Panel: Country specific factors. Right Panel: Country and US factors.

Moving on to the U.K., we find the FTSE excess returns to be predictable primary by the jump intensity, but there is weak evidence of some predictability by the variance factor in this case. However, when we include the U.S. factors into the excess return regressions, the U.S. left jump intensity enters as the only strong predictor, with an impressive increase in overall R^2 from below 20% at the half year horizon to well above 30%. In this case, the pricing of the U.S. tail factor emerges as the primary forecast tool for the British equity risk premium.

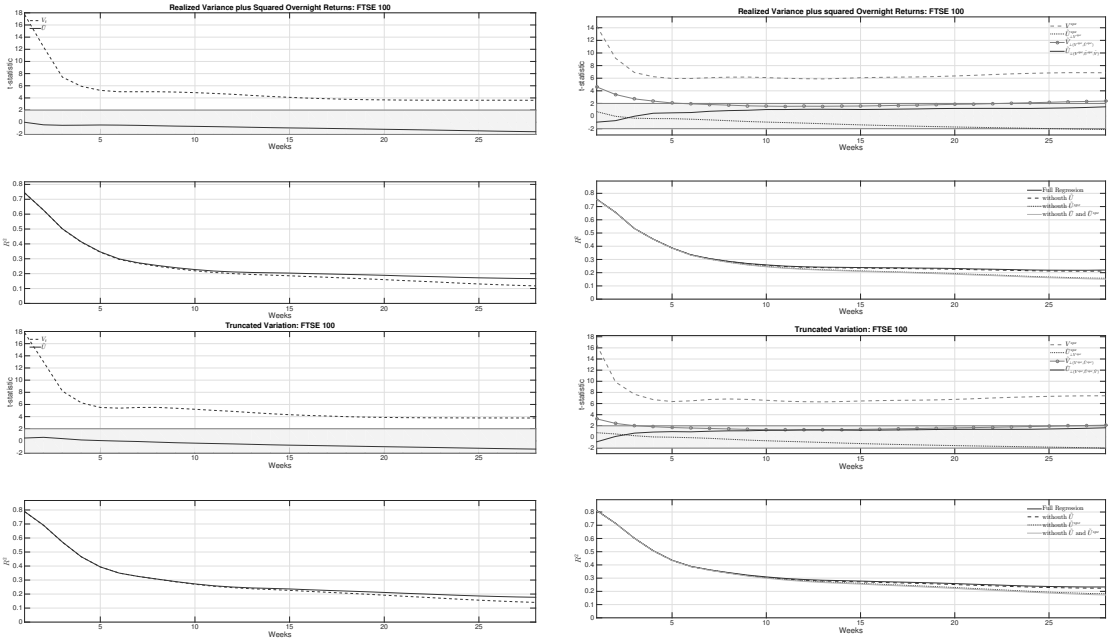

Figure 11: **Predictive regression, FTSE Return Variation Measures**. Top Panels: RV plus squared overnight returns. Bottom Panels: Truncated Variation. Left Panels: country specific factors. Right Panels: country specific and US factors.

For the U.K. return variation regressions, we recover the usual result: the spot variance factors in the two countries are highly correlated and both predict the future return variation very well. The incremental explanatory power of the U.S. variance factor is very limited.

Finally, as an example of a country that was deeply affected by the European sovereign debt crisis, we consider the predictive regressions for the Spanish IBEX index. The equity risk premium is again largely predictable from the jump intensity factor, but the statistical significance is weaker than in the prior cases. Notably, when including the U.S. tail factor in the regression, we find the explanatory power to increase significantly and the incremental information of the local Spanish jump factor dominates. This suggests the presence of separate global and local risk premium indicators, and accounting for their joint effect allows for better identification of the local Spanish component. The heightened importance of local risk premium factors in countries confronting unique sovereign debt issues is, of course, also economically intuitive.

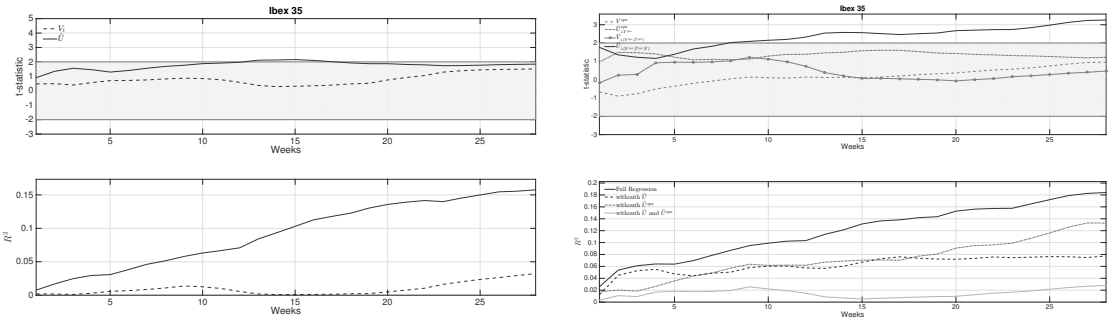

Figure 12: **Predictive Regressions, IBEX Excess Returns**. Left Panel: Country specific factors. Right Panel: Country plus US factors.

For the Spanish return variation, we recover the usual irrelevance of the jump intensity factors, while the spot variance factors are highly significant. The only noteworthy discrepancy relative to our prior evidence is a comparatively stronger significance of the local volatility factor. Again, this is consistent with the idea that local economic conditions are relatively more important in Spain than in countries that were less impacted by the European crisis.

Overall, we conclude that there is a clear separation in the determinants of equity risk premium and market risks: the latter is well captured by the current level of market volatility while the former is driven by the option-implied risk-neutral jump intensity that is not spanned linearly by the market volatility. Our evidence suggests that this phenomenon is in force not only for US, but also for the other markets in our analysis. This finding can also explain the international evidence for the predictive power of the country-specific variance risk premium for that country's future

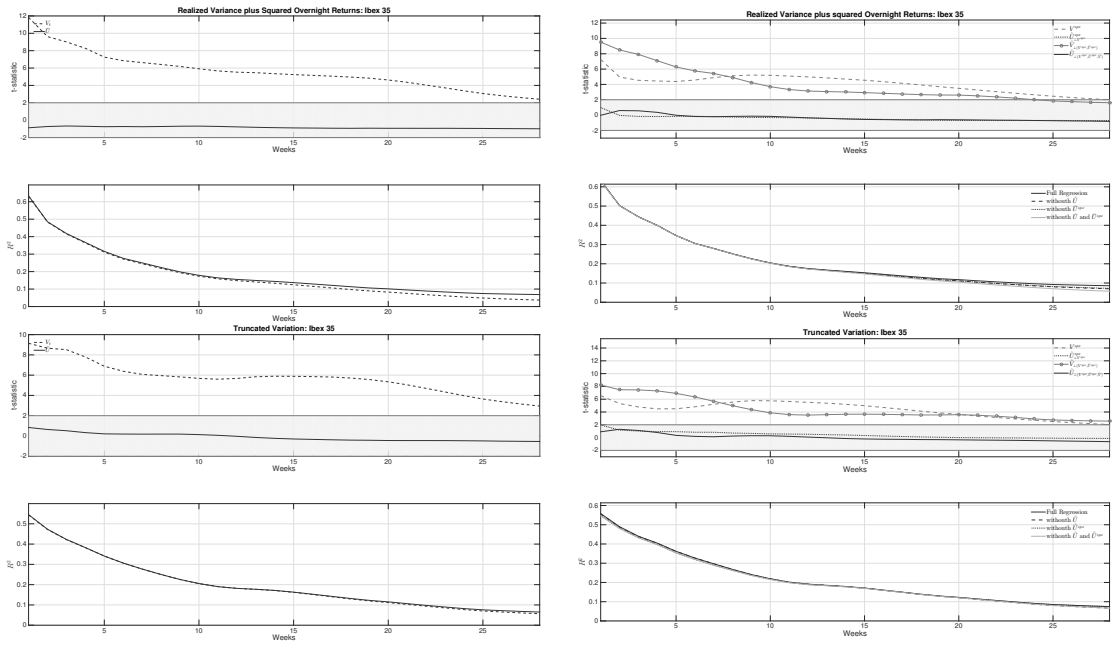

Figure 13: **Predictive Regression, IBEX Return Variation**. Top Panels: RV plus squared overnight returns. Bottom Panels: Truncated Variation. Left Panels: Country specific factors. Right Panels: Country and US factors.

excess returns documented in Bollerslev et al. (2014). The variance risk premium represents the difference between priced and expected future volatility, and a large piece of this differential stems from the compensation for negative jump risk which, in turn, is driven by the risk-neutral jump intensity factor U, that we uncover here.[6]

6.4 Implications for the Negative Jump Risk Premium

We now study the implications of the option estimates for the risk premia due to the negative jumps, i.e., we are interested in

$$\begin{aligned} NJP_{t,t+h} &= \frac{1}{h}\mathbb{E}_t^{\mathbb{Q}}\left(\sum_{s\in[t,t+h]}(\Delta X_s)^2 1_{\{\Delta X_s<0\}}\right) - \frac{1}{h}\mathbb{E}_t^{\mathbb{P}}\left(\sum_{s\in[t,t+h]}(\Delta X_s)^2 1_{\{\Delta X_s<0\}}\right) \\ &= \frac{1}{h}\mathbb{E}_t^{\mathbb{Q}}\left(\int_t^{t+h}\int_{x<0} x^2\nu_s^{\mathbb{Q}}(dx)ds\right) - \frac{1}{h}\mathbb{E}_t^{\mathbb{P}}\left(\int_t^{t+h}\int_{x<0} x^2\nu_s^{\mathbb{P}}(dx)ds\right). \end{aligned} \tag{8}$$

The risk-neutral expectation in $NJP_{t,t+h}$ can be easily computed from the estimates of our risk-neutral option pricing model. On the other hand, the $\mathbb{P}$ expectation of $NJP_{t,t+h}$ can be computed

[6]Our findings here are also related with earlier large literature documenting return predictability on international stock market level, see e.g., Harvey (1991), Bekaert and Hodrick (1992), Campbell and Hamao (1992), Ferson and Harvey (1993) and Hjalmarsson (2010).

from the linear projection of the realized jump variation on the state vector that summarizes the relevant information for forecasting at a given point in time. The results for a horizon of four months are presented on Figure 14. The negative jump risk premium peaks in the periods of the financial crisis and the two subsequent European sovereign debt crises and then gradually subsides after these events. It is useful to contrast the dynamics of $NJP_{t,t+h}$ with that of the risk-neutral spot jump variation displayed on Figure 5. Recall that the latter is, simply, $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx)$, and it thus contains both the expected future negative jump risk (i.e., its $\mathbb{P}$ counterpart) as well as the compensation for it. The effect of the three crises on the risk premium is far more balanced than the corresponding effect on the total risk-neutral negative jump variation shown in Figure 5. This means that the expensiveness of the short-maturity out-of-the-money put options in relative terms, during and immediately after the financial crisis of 2008, was partially due to the expected elevation in the level of risk at the time. Once this is accounted for in $NJP_{t,t+h}$, the difference in the latter across the crises periods becomes much smaller. We note that our risk premium series bear resemblance to the pure jump intensity factor, i.e., the component of U that is not in the linear span of V; compare, e.g., to the bottom panel of Figure 1 for the US evidence.

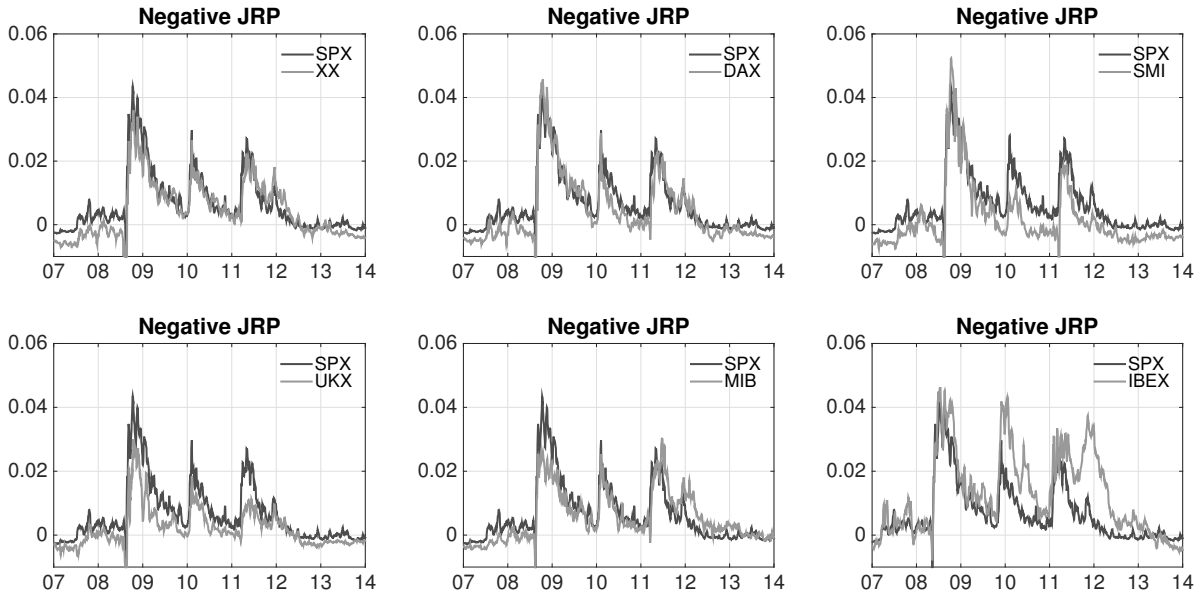

Figure 14: **Negative Jump Risk Premium**.

Moving on to the risk premiums inferred for the different countries, we notice some interesting patterns. First, there is a lot of similarity in the evolution of the premiums across US, UK, Germany and Switzerland. Similar to the risk-neutral negative jump variation comparison, the risk premiums

in UK and Switzerland is somewhat smaller than the the one in US, but the pattern of the dynamics is the same across these countries. Second, the negative jump risk premiums in Italy and Spain show sizable deviations from the benchmark in the US. The risk premium in Italy is somewhat smaller than the one in the US during the financial crisis in 2008, it is largely identical during the first sovereign debt crisis of 2010, and it exceeds that in US during and after the second sovereign debt crisis. On the other hand, the negative jump risk premium in Spain tends to be higher than the one in US during most of the sample. The "fear" of a downturn in Spain surges well above the level in the US during the first sovereign debt crisis and then gradually subsides to the level of the US. This behavior is repeated following the second sovereign debt crisis. Seemingly, the option market in Spain reflects much deeper concerns about a possible downturn in the economy during the European debt crises than the Italian option market does for that economy economy. This behavior is interesting, particularly because the actual jump risks in the two countries turn out to be very similar. To illustrate this, on Figure 15, we plot a scatter of the detected jumps in the two markets over our full sample period. At the times of the common jumps, the two markets react very similarly. Indeed, a regression of the jumps in Spain against those in Italy, at the common jump times, yields a slope coefficient of 0.99 with an R^2 of 0.86. Furthermore, the jump variation of the disjoint jumps, i.e., ones that are present only on the Italian or Spanish market, are very similar. In spite of this striking similarity in the realized jump risk, the data reveal that the negative jump risk premiums across the two markets behave differently.

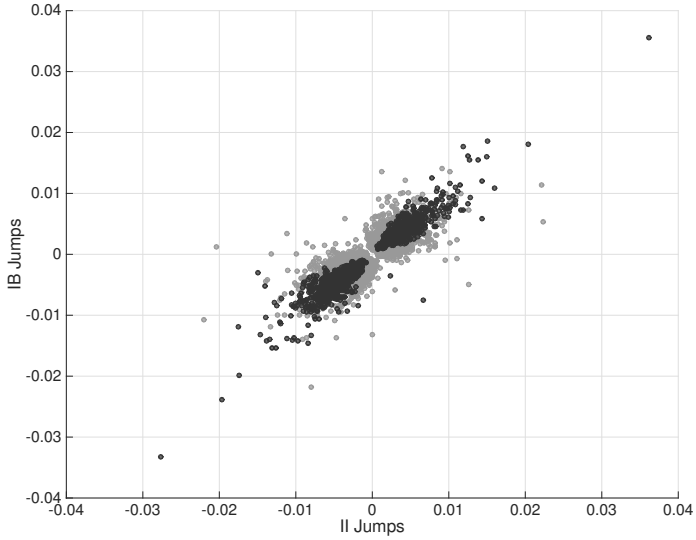

Figure 15: **Scatter of Market Jumps in Spain vs. Italy**. Black dots correspond to commonly arriving jumps and light-grey ones to the ones that are disjoint.

7 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. The results confirm robustness of the original empirical findings for the US market. 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. On the other hand, the spot variance does not generally predict the equity risk premium, and the tail factor has no incremental explanatory power for the future return variation. Thus, the option surfaces afford us with a split into separate risk and risk premium (or risk pricing) factors. Importantly, standard approaches exploiting only volatility factors will 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 very 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. The details are intriguing, with Germany only displaying 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 initial sovereign debt crisis. In contrast, the Spanish tail factor is exceedingly elevated and only return to more normal levels years after the crises. As such, the Spanish risk pricing is clearly unique within the set of indices explored.

We conclude that the equity-index options contain critical information about the expected future evolution of financial risk as well as the pricing of the relevant risk factors. Moreover, the option markets are integrated with the underlying equity markets so that the state of the option surface provides important insights into the size of likely evolution of the equity and variance risk premiums. This should set the stage for more detailed work exploring the equity risk pricing using joint information from the equity and the derivatives markets.

Given our specific findings regarding the relative risk pricing in Italy and Spain, it would also be of interest, in future work, to associate the inferred downside tail risk premium in Spain with the actual sequence of economic events affecting the country during this period.

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). Pricing Short-Term Market Risk: Evidence from Weekly Options. Working paper, Northwestern University.

Andersen, T. G., N. Fusari, and V. Todorov (2015c). The Risk Premia Embedded in Index Options. *Journal of Financial Economics 117*, 558–584.

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. and V. Todorov (2014). Time- Varying Jump Tails. *Journal of Econometrics 183*, 168–180.

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.

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.

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.

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

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

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

Li, J. and G. Zinna (2015). The Variance Risk Premium: Components, Term Structures, and Stock Return Predictability. Working paper.

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.

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.

8 Appendix

8.1 Options Data

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`.

8.2 Daily Returns Data

8.3 Construction of High-Frequency Measures

We download the data using the TickWrite software provided by TickData. We select Time Based Bars interval with one minute granularity holding the last value in case there is no price over that particular minute. We also consider (for Futures) the front contract and we use the Auto Roll method provided by the software. We only consider the daily section combining pit and electronic trading.

After this first stage we proceed with the cleaning of the data. Specifically we apply the following filters:

	ES	XX	DA	SW	FT	II	IB
ES	1.00	0.90	0.89	0.65	0.88	0.59	0.59
XX		1.00	0.95	0.72	0.89	0.75	0.76
DA			1.00	0.69	0.87	0.68	0.67
SW				1.00	0.71	0.73	0.73
FT					1.00	0.64	0.63
II						1.00	0.87
IB							1.00

Table 5: **Index returns correlation**.

1. we keep only observations between Monday and Friday.
2. we remove days with no price changes and days corresponding to US holidays
3. we remove days with number of observations less than the average number of daily observations (for example for the SP500 we have usually 405 one-minute observations). This filter removes half-trading days such has (in case of the US) the day before Thanksgiving or the day before Christmas.

With the cleaned data we then construct the high-frequency measures following these steps:

1. we create a grid with 5 minutes returns.
2. Jumps Detection:
   - We compute, for each series, the Time of the Day (TOD) function. We recompute the TOD each time the exchange changed the trading hours.
   - at each point in time, starting on the second day, we compute the RV (Realized Variance) and the BV (Bipower Variation) over the previous 24 hours and we use them to detect a jump if the subsequent log-return:

$$|\Delta_i^{n,t} f| > 4(RV_t \wedge BV_t)\Delta_{n,t}^{0.49} * TOD_i.$$

   - After we detect the jumps we construct the high-frequency measures

- For the computation of the Local Truncated Variation we take into account for the following:
  - we use the last 3 hours before the closing of the option market;
  - 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;
  - We use the same rational in computing the integrated quarticity (IQ).

8.4 Parameter Estimates

8.4.1 SPX

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

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

Table 6: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day XX options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.2 EUROSTOXX

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

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

Table 7: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day SX5E options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.3 DAX

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

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

Table 8: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day DAX options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.4 SWI

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

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

Table 9: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day SMI options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.5 UKY

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

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

Table 10: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day UKX options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.6 MIB

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

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

Table 11: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day MIB options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.4.7 IBEX

Panel A: Parameter Estimates								
Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ	-1.000	0.047	κ_u	1.285	0.146	μ_v	4.397	0.620
$\overline{v}$	0.040	0.001	c_0^+	6.382	2.349	μ_u	55.010	20.512
κ_v	3.514	0.121	λ^-	13.603	0.614	ρ_j	0.000	0.000
σ_v	0.529	0.027	λ^+	40.044	5.797			

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

Table 12: **Estimation results for the parametric model defined by equation** (1)**.** The results are based on end of day IBEX options, covering January 2007 – December 2014. **Panel A** reports the parameter estimates obtained using weekly observations on Wednesday, or the closed business day in case of a market closure on Wednesday. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

8.5 Alternative Two-Factor Volatility Model

The alternative two-factor model that we compare our model in equation (1) is given by

$$\begin{aligned}
\frac{dX_t}{X_{t-}} &= (r_t - \delta_t)\,dt \;+\; \sqrt{V_{1,t}}\,dW^{\mathbb{Q}}_{1,t} \;+\; \sqrt{V_{2,t}}\,dW^{\mathbb{Q}}_{2,t} \;+\; \int_{\mathbb{R}^2} (e^x - 1)\,\widetilde{\mu}^{\mathbb{Q}}(dt,dx)\,, \\
dV_{1,t} &= \kappa_1\,(\overline{v}_1 \,-\, V_{1,t})\,dt \;+\; \sigma_1\,\sqrt{V_{1,t}}\,dB^{\mathbb{Q}}_{1,t} \;+\; \mu_1 \int_{\mathbb{R}^2} x^2\,1_{\{x<0\}}\,\mu(dt,dx)\,, \\
dV_{2,t} &= \kappa_2\,(\overline{v}_2 \,-\, V_{2,t})\,dt \;+\; \sigma_2\,\sqrt{V_{2,t}}\,dB^{\mathbb{Q}}_{2,t}\,,
\end{aligned} \tag{9}$$

where $(W^{\mathbb{Q}}_{1,t}, W^{\mathbb{Q}}_{2,t}, B^{\mathbb{Q}}_{1,t}, B^{\mathbb{Q}}_{2,t})$ is a four-dimensional Brownian motion with $\operatorname{corr}\left(W^{\mathbb{Q}}_{1,t}, B^{\mathbb{Q}}_{1,t}\right) = \rho_1$ and $\operatorname{corr}\left(W^{\mathbb{Q}}_{2,t}, B^{\mathbb{Q}}_{2,t}\right) = \rho_2$. In addition, μ is an integer-valued measure counting the jumps in the price, X. The corresponding (instantaneous) jump intensity, under the risk-neutral measure, is $dt \otimes \nu^{\mathbb{Q}}_t(dx)$. The difference, $\widetilde{\mu}^{\mathbb{Q}}(dt,dx) = \mu(dt,dx) - dt\,\nu^{\mathbb{Q}}_t(dx)$, 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, V_1 displays (positive) jumps proportional to x^2 with proportionality factor given by μ_1. We model the jump intensity as:

$$\frac{\nu^{\mathbb{Q}}_t(dx)}{dx} \;=\; c^-(t)\cdot 1_{\{x<0\}}\,\lambda_-\,e^{-\lambda_-|x|} \;+\; c^+(t)\cdot 1_{\{x>0\}}\,\lambda_+\,e^{-\lambda_+ x}\,. \tag{10}$$

The two terms on the right hand side refer to negative and positive price jumps, respectively. We model the price jumps as exponentially distributed, with separate tail decay parameters, λ_- and λ_+, for negative and positive jumps. Finally, the time-varying jump intensities are governed by the $c^-(t)$ and $c^+(t)$ coefficients. These coefficients evolve as affine functions of the state vector,

$$c^-(t) \;=\; c^-_0 + c^-_1 V_{1,t-} + c^-_2 V_{2,t-}, \qquad c^+(t) \;=\; c^+_0 + c^+_1 V_{1,t-} + c^+_2 V_{2,t-}\,. \tag{11}$$

This representation involves a large set of parameters that can be hard to identify separately. At the estimation stage, we zero out various insignificant coefficients. Specifically, for the negative jump intensity, we fix c^-_0 to zero. Also, we set c^+_1 and c^+_2 to zero, implying that the intensity of positive jumps is constant, $c^+(t) = c^+_0$.