

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

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

August 2018

Abstract

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

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

JEL classification: G12, G13, G15, G17.

*Andersen gratefully acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The work is partially supported by NSF Grant SES-1530748. We thank the Editor, Todd Clark, the Associate Editor, and two anonymous referees for many useful comments and suggestions. We thank Eben Lazarus for providing the code used to implement the HAR estimators in Lazarus et al. (2018). We also thank participants at the following presentations for useful comments: AQR Capital Management; Midwest Finance Association Meeting 2016, Chicago; Key-note address at the 2016 Nippon Finance Association Meeting, Yokohama; 30th Anniversary of GARCH Models and Measures Conference, Toulouse School of Economics; Key-note address at the Financial Econometrics and Empirical Asset Pricing Conference 2016, Lancaster University Management School; SoFiE Annual Meeting 2016, City University, Hong Kong; NBER Summer Institute 2016; Financial Econometrics Conference 2016, Questrom School of Business, Boston University; CBOE Conference on Volatility and Derivatives 2016, Chicago; Financial Econometrics Conference 2016, Rio de Janeiro; as well as participants at numerous seminars.

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

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

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

1 Introduction

Over the last decade, global financial markets have been roiled by major shocks, including the financial crisis originating in October 2008 and European debt crises in May 2010 and August 2011. These events represent a challenge to dynamic asset pricing models. Can they accommodate the observed interdependencies between volatility and tail events, and their pricing? How do volatility and tails differ across countries? How do the variance and tail risk premiums relate to the compensation for exposure to equity risk, i.e., the equity premium? The increasing liquidity of derivative markets worldwide provides an opportunity to shed light on these questions. In particular, the trading of equity-index options has grown sharply, with both more strikes per maturity and additional maturities on offer. Most importantly, there has been a dramatic increase in the trading of options with short tenor. As a result, we now have access to active prices for securities that embed rich information about the pricing of volatility and tail risk in many separate countries. By combining the pricing of financial risks, embodied within the equity-index option surfaces, with ex-post information on realized returns, volatilities and jumps, we can gauge the relative size of the risk premiums and what factors drive the risk compensation.

In the current work, we draw on daily observations for option indices in the U.S. (S&P 500), Euro-zone (ESTOXX), Germany (DAX), Switzerland (SMI), U.K. (FTSE), Italy (MIB), and Spain (IBEX) over 2007-2014 to extract factors that are critical for the pricing of equity market risk across North America and Europe. The performance of the individual equity indices varies drastically, with the German market appreciating by an average of about 8% per year and the Italian depreciating by 5% annually, as illustrated in Figure 1. This heterogeneity provides additional power in studying the connection between market tail risk and the equity risk premium.

Most existing option pricing models capture the dynamics of the equity-index option surface through the evolution of factors that determine the volatility of the underlying stock market, see, e.g., Bates (1996, 2000, 2003), Pan (2002), Eraker (2004) and Broadie et al. (2007). However, recent evidence suggests that the fluctuations in the left tail of the risk-neutral density, extracted from equity-index options, can be spanned neither by regular volatility factors nor by realized risk measures constructed from return data. Hence, a distinct factor is necessary to account for the priced downside risk in the option surface, see Andersen et al. (2015b). We accommodate such features by specifying a parametric risk-neutral return model, but imposing no structural




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

assumptions on the underlying return generating process. Other models allowing for an element of separation between jump and volatility factors include Santa-Clara and Yan (2010), Christoffersen et al. (2012), Gruber et al. (2015), and Li and Zinna (2018).

We seek to fully exploit the available information by utilizing *all* option quotes that satisfy our mild filtering criteria. This is important because the European samples are short and display varying liquidity. Such features render day-by-day factor identification more challenging than for the U.S. market. Consequently, we specify a parsimonious risk-neutral model with a single volatility component along with the tail factor. The short-dated options are highly informative regarding the current level of volatility and jump intensities, while the longer-dated options primarily speak to the factor dynamics. In combination, this facilitates robust identification of the volatility and jump components and reduces the number of parameters in the risk-neutral model, providing a solid basis for out-of-sample exploration of the predictive power of the option-implied factors.

We briefly summarize our empirical findings. Most importantly, we find a substantial discrepancy between the time series evolution of priced tail risk and volatility for every option market we analyze. A common feature emerges in the aftermath of crises – the left tail factor is correlated with volatility, yet it remains elevated long after volatility subsides to pre-crisis levels. This is incompatible with the usual approach to the modeling of volatility and jump risk in the literature.

The stark separation of the tail and volatility factors has important implications for the pricing

and dynamics of market risk. Even though volatility is a strong predictor of future market risk, it provides no forecast power for the equity risk premium. In contrast, our left tail factor does not help predict future market risks, yet it has highly significant explanatory power for future market returns. The tail factor also constitutes the primary driver of the negative tail risk premium, suggesting this is the operative channel through which it forecasts the equity premium. In particular, following crises, equity prices are heavily discounted and the option-implied tail factor remains elevated, even as market volatility resides. Throughout our sample, this combination serves a strong signal that the market will outperform in the future.

These findings inspire a number of new, interrelated hypotheses regarding the relative forecast power of option-implied and return variation measures for the equity risk premium, which we validate for each of our equity indices. For example, the compensation for exposure to return variation, i.e., the variance risk premium or VRP, has been intensely studied following the work of Carr and Wu (2009) and Bollerslev et al. (2009). We find, most strikingly, that the VRP only forecasts future equity returns because it embeds the left tail factor. Once the jump risk premium is stripped from the VRP, it no longer provides a signal for the performance of the equity market. Since the VRP constructed from different asset classes also have predictive power for the corresponding excess returns, e.g., credit spreads (Wang et al. (2013)), currency returns (Londono and Zhou (2017)), and bond returns (Mueller et al. (2017)), our results likely have broader implications for risk pricing across diverse global security markets.

In terms of the volatility factor, we find the European and U.S. markets to evolve in near unison through the financial crisis of 2008-2009. In contrast, we observe discrepancies during and after the initial European sovereign debt crisis. Overall, the U.K., Swiss and, to some extent, German volatilities remain close throughout the sample. The largest divergences in volatility dynamics occur between the U.S., U.K. and Swiss (non euro-zone) indices on one side and the Spanish and Italian on the other, with the latter representing the Southern euro-zone countries in our sample.

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

Exploiting the extracted factor realizations along with estimates for the risk-neutral and actual

return dynamics, we further document very strong dependence between the negative jump risk premiums for the equity indices throughout the sample. By contrast, the dependence between the diffusive risk premiums across markets is strong only for the financial crisis period 2007-2009, while for 2010-2014, associated with the European sovereign debt crises, the dependence between the diffusive risk premiums weakens significantly. Given the critical role of the tail factor for future equity returns, the coherence in the compensation for tail risk across indices rationalizes the strong correlation of the drift component among the individual equity indices, evident in Figure 1, despite the heterogeneity in the severity of the shocks affecting the individual countries.

Our work is related to an evolving literature on the ability of equity option and equity return variation measures to capture the pricing of exposure to volatility and tail, jump or skew risk. By extension, such measures should possess forecast power for future equity returns. First, there is old literature studying empirically the mean variance tradeoff and the ability of stock market volatility to predict future returns, see e.g., French et al. (1987) and Glosten et al. (1993) for early references. More recently, Almeida et al. (2017), Bali et al. (2009) and Kelly and Jiang (2014) use different approaches to construct tail measures from the cross-section of observed stock prices and show that these measures can predict future market returns.

The above cited work does not use derivatives in the construction of the tail and volatility measures. Employing individual equity options, Bali and Hovakimian (2009) find that both the variance risk premium and an alternative option spread measure—interpreted as a jump risk measure—explain cross-sectional variation in expected returns. Later, Conrad et al. (2013) find that more negative risk-neutral volatility and skew measures are associated with higher returns in the cross-section. The latter is consistent with the evidence of preferences for positive skew equity assets obtained in Bali and Murray (2013).

For the U.S. equity index, Bollerslev and Todorov (2011) document that a nonparametric-based left tail premium measure explains a significant fraction of both the variance and equity risk premium. Bollerslev et al. (2015) and Du and Kapadia (2012) show that nonparametric jump tail measures constructed from options can predict returns in the future over horizons of several months, while Atilgan et al. (2013) find evidence for short-term return predictability (daily and weekly) by option-implied measures. Meanwhile, Kozhan et al. (2013) demonstrate that the skew risk premium is substantial, but overlaps with the variance risk premium, and conclude that these

risk factors have a common source. Finally, Li and Zinna (2018) find that both the level and slope of the term structure for the variance risk premium provide significant predictors of future equity returns.

On the international front, Bollerslev et al. (2014) show that local VRPs provide qualitatively similar, yet decidedly weaker, predictive power for national equity-index returns than in the U.S. However, a weighted average of the VRP measures provides a global proxy with significantly improved explanatory power across the indices. Similarly, Londono (2016) find that foreign VRP measures are insignificant predictors for the local equity returns, but the U.S. VRP provide much superior performance, with only the Japanese index remaining unpredictable. Finally, Gao et al. (2018) establish predictive power for the excess returns on a wide set of asset classes using a global tail index, constructed in accordance with the “jump and tail” index of Du and Kapadia (2012).

Our findings are consistent with existing results in many respects. The VRP has predictive power for equity returns, but negative tail measures perform even better; the pricing of equity risk has a strong global component; and equity-index returns are highly correlated across countries.

Yet, our contribution deviates from the prior literature on numerous points. Most importantly, we identify a component in the risk-neutral left jump intensity, orthogonal to spot volatility, which we show is the main driver of the equity return predictability. Raw option-implied risk measures (volatility or tail related) have weaker predictive ability, rationalizing some of the existing empirical results in the literature. Our findings suggest the existence of an important wedge between risk and risk premium dynamics for all the equity indices in our analysis.

Two, in contrast to most prior studies, we exploit the information across all option maturities and strikes. Earlier work typically relies on one maturity only and aggregate many options into a single measure, such as the volatility level or skew. Our use of short-dated options enhances our ability to identify the current state of the volatility and jump factors. Moreover, the observations for the longer tenors facilitate identification of the risk-neutral volatility and jump dynamics.

Three, we decompose the variance risk premium into a number of separate elements and provide a strict rank ordering of their ability to forecast equity-index returns. In particular, the continuous return variation and associated premium should have no predictive power. These new hypotheses are corroborated for each of the equity indices.

Four, we obtain strong empirical results for return predictability based on the foreign option

panels. The U.S. evidence remains highly significant, but not more so than for many of the other countries. This finding is consistent with our approach affording an improved extraction of information from the smaller sample of options available outside the U.S.

Finally, the only existing work we are aware of linking tail measures with equity return predictability internationally is Gao et al. (2018). Their tail index captures the variation in the higher order risk-neutral moments induced by jumps, which is controlled by the variation in the jump intensity. In contrast to our left tail factor, the tail index of Gao et al. (2018) covaries quite strongly with volatility. Moreover, by construction, their index puts most of the weight on very deep out-of-the-money option prices. This may render the tail measure somewhat unstable, as the active strike coverage is subject to fluctuations in liquidity and the contracts are subject to relatively large measurement errors. By contrast, we are less sensitive to liquidity and data availability issues, but at the cost of imposing some parametric structure on the option dynamics.

The paper is organized as follows. Section 2 presents the model used to fit the option surfaces, and Section 3 reviews the estimation method. Section 4 focuses on the option-implied factors and their predictive ability. Section 5 explores the interaction among the factors and the risk premiums, with emphasis on the role of the jump factor. We study the common features in the option-implied factors and risk premiums in Section 6. Finally, Section 7 concludes. Additional details on data, estimation method and results are provided in the online Appendix.

2 Model

This section introduces our model for the risk-neutral dynamics of a generic equity market index, denoted X , and defines a set of related risk measures and premiums explored in the empirical work.

The model is designed to capture the dynamics of the option surface through a low-dimensional state vector within a continuous-time no-arbitrage setting. We subsequently use the option-implied state vector to predict future equity risks and risk premiums and relate these to sources of priced risk in the economy. Because the European option samples are limited relative to the U.S., we seek a robust and parsimonious specification that captures the salient features and facilitates a sharp day-by-day separation of the volatility and jump factors. Towards that end, we are inspired by prior successful representations, but limit ourselves to one volatility and one jump intensity factor.

2.1 The Risk-Neutral Equity-Index Dynamics

Our two-factor model for the risk-neutral index dynamics is given by the following restricted version of the representation in Andersen et al. (2015b),

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

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

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

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

Following Kou (2002), we assume that the price jumps are exponential, with separate tail decay parameters, λ_- and λ_+ , for negative and positive jumps. Finally, the left jump intensity is governed by the factor U_t while the positive jump intensity is constant and equal to c_0^+ . The latter assumption is innocuous. We verify that all qualitative empirical conclusions are robust to standard specifications of time variation in the positive jump factor.

Our jump modeling includes novel features and deviates markedly from the standard parametric specification in the empirical option pricing literature. First, the price jumps are exponentially distributed, and consequently have fat tails, while most prior studies rely on Gaussian price jumps,

following Merton (1976). Second, the jumps in the factors V and U are linked deterministically to the negative price jumps, with the squared jumps impacting the factor dynamics in a manner reminiscent of discrete GARCH models. Third, and most importantly for the analysis here, the jump intensity is time-varying, and with innovations that are linked directly to the volatility and price jumps, yet its dynamics is largely decoupled from volatility. This is unlike most existing models, with the notable exception of Christoffersen et al. (2012), Gruber et al. (2015) and Li and Zinna (2018). Nonetheless, the representation (1) belongs to the affine class of models of Duffie et al. (2000).¹ For future reference, we label our two-factor affine model the $2FU$ model.

The model is motivated by recent evidence that the negative jump intensity represents a separate risk factor which is essential in capturing the dynamics in the left side of the option surface. Moreover, that study finds the left tail factor to be a critical driver of the variation in the equity and variance risk premiums. At the same time, model (1) provides a more parsimonious representation than the preferred specification in Andersen et al. (2015b), which features two separate volatility factors. This choice reflects our desire to achieve precise identification of the relevant jump factor from the shorter, less liquid option samples encountered in this study. As discussed in Section 4 below, the jump factor is robustly identified in the current more restricted model.

2.2 Risk Measures and Risk Premiums

We conclude the section by introducing a few risk measures and associated risk premiums that are used throughout the paper. The quadratic variation of the log-price is given by,

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

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

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

where $CV_{t,t+h}$ and $JV_{t,t+h}$ are the expected diffusive and jump variation under the risk-neutral measure, respectively.² The expected jump variation $JV_{t,t+h}$, can be decomposed further into

¹There is some option-based evidence, see, e.g., Jones (2003) and Christoffersen et al. (2010), that non-affine models work better. The fact that our key findings are driven primarily by short-maturity options should mitigate the importance of potential misspecification stemming from nonlinearities in the volatility dynamics.

²For notational simplicity, we often do not include a superscript \mathbb{Q} to indicate that a specific quantity is obtained under the risk-neutral measure. For example, the expected diffusive return variation under the risk-neutral measure

terms stemming from negative and positive jumps,

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

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

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

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

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

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

and the negative jump risk premium as,

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

Our computation of the expected return variation measures under the statistical (\mathbb{P}) measure follows standard high-frequency procedures for forecasting the volatility and jump risks. That is, we do not specify a parametric model for X under \mathbb{P} . Instead, we use a reduced-form model to generate the necessary conditional \mathbb{P} -expectations for the calculation of the requisite risk premiums, exploiting the popular HAR model of Corsi (2009). Consistent with our empirical evidence in Section 4.2.1 below, we find only lagged continuous return variation measures, and not lagged jump variation measures, to be significant predictors of the future return variation. Appendix A.3 provides a detailed description of these procedures.

is denoted $CV_{t,t+h}$ rather than $CV_{t,t+h}^{\mathbb{Q}}$. This should not cause confusion as, whenever we refer to expectations under the actual or statistical probability measure, the relevant quantities carry the superscript \mathbb{P} .

3 Estimation Procedure

We follow the inference procedures in Andersen et al. (2015a) for recovering the parameters and latent factor realizations of model (1). The option prices are converted into Black-Scholes implied volatilities (BSIV), i.e., any out-of-the-money (OTM) option price observed at time t with tenor τ 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 parameter vector θ , the model-implied BSIV is denoted $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$. Estimation now proceeds by minimizing the distance between observed and model-implied BSIV in a metric that also penalizes for discrepancies between the inferred spot volatilities and those estimated (in a model-free way) from high-frequency data on the underlying asset, $\sqrt{\widehat{V}_t^n}$, see Appendix A.3 for further details. The imposition of (statistical) equality between the spot volatility estimated from the actual and risk-neutral measure reflects a basic no-arbitrage condition implied by the underlying option pricing paradigm.

To formally specify the estimation criterion, we introduce additional notation. Let $t = 1, \dots, T$, denote the dates for which we observe the option prices at the end of trading. We focus on OTM options, with $k \leq 0$ indicating OTM puts and $k > 0$ OTM calls. Due to put-call parity and the lower liquidity of in-the-money options, there is little loss of information from using only OTM options for estimation.

We obtain point estimates for the risk-neutral parameter vector θ and period-by-period state vector $\mathbf{S}_t = (V_t, U_t)$ from the following optimization problem, see Andersen et al. (2015a),

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

To reduce the computational burden, we estimate the system for options only sampled on Wednesday or, if this date is missing, the following trading day. Credible identification of the system is obtained by observing heterogeneous constellations of the option surface across time. We facilitate this objective by exploiting all, including very short-dated, options that satisfy a mild filter, designed to identify whether a given option observation is flawed. The surface varies dramatically for the early and late years relative to the periods around the financial and European debt crises. Once the parameter vector and the state variable realizations for all Wednesdays have been obtained, it is straightforward to “filter” the state variables for the remaining trading days, exploiting the

estimated parameters, option data, and the criterion (6). Thus, we obtain daily estimates for the state realizations, even if full-fledged estimation is performed only for weekly data.

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

4 Option Factors

The data used for our empirical analysis of the U.S., Eurozone, German, British, Swiss, Italian, and Spanish equity indices are described in Appendix A.1. As discussed in Section 2, our model is restricted relative to Andersen et al. (2015b). We confirm, however, that the analysis is robust to a more refined modeling of the volatility dynamics by documenting near identical jump factor extraction from the alternative models in Appendix A.5. Additional details on our estimation results are provided in Appendix A.6.

4.1 Country-by-Country Factor Realizations

This section explores the implied factor realizations obtained from model (1) and the estimation procedure outlined in Section 3. We first compare the spot variance of the European indices, benchmarking against the S&P 500 results.

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

In contrast, notable discrepancies between the S&P 500 and German DAX emerge during the second phase of the European sovereign debt crisis, when the DAX volatility spikes more, and a positive gap remains from that point onward, albeit to a varying degree. For the broader ESTOXX

index, the same effect is clearly visible and originates with the initial phase of the debt crisis. Thus, while the volatility patterns were similar for all the indices through the financial crisis, the response to the sovereign debt crisis is heterogeneous, with the impact reflecting prior perceptions regarding the sensitivity of the respective economies to the European crisis. This is especially striking for the Italian and Spanish indices, as both react strongly to the crisis events, but with varying amplitudes across the main episodes. These Southern European indices reach a volatility plateau well above the others ever since the sovereign crises surfaced in early 2010. This systematic divergence lowers the volatility correlations for MIB and IBEX with the S&P 500 to 0.75 and 0.77, respectively.

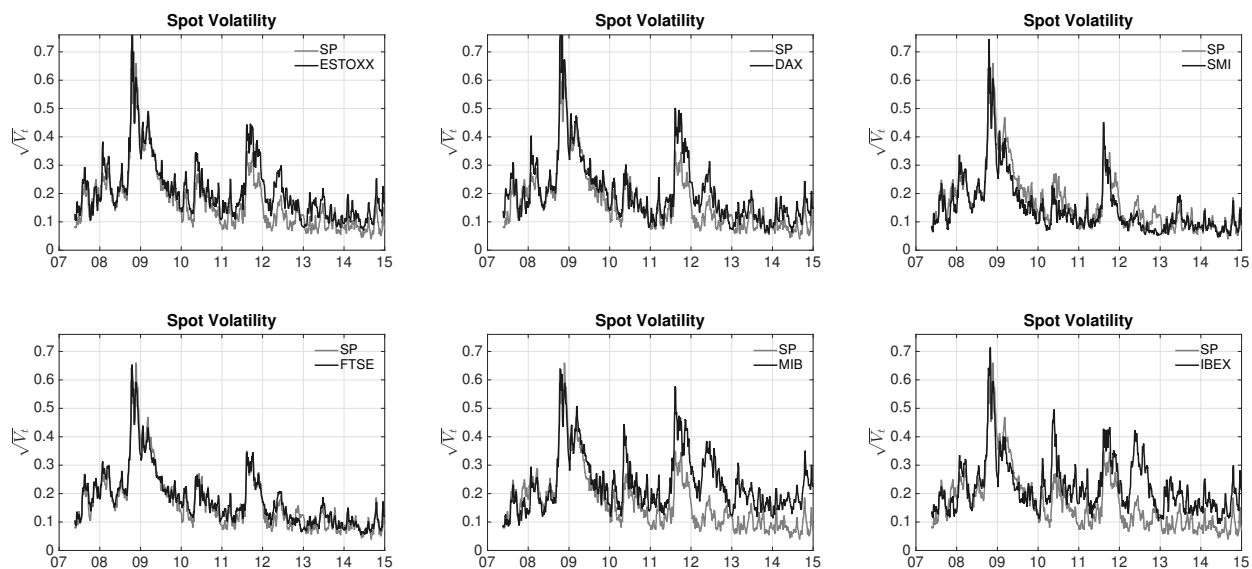


Figure 2: **Volatility Factor Comparison.** For each option-implied spot variance factor, obtained at the close of the trading day, we report the trailing five-day moving average of $\sqrt{V_t}$. The series are given in decimals and refer to annualized values.

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

At first glance, the pattern is similar to the one observed for the volatility factors. This is natural as the volatility and jump factors are highly correlated for all countries. Nevertheless, the

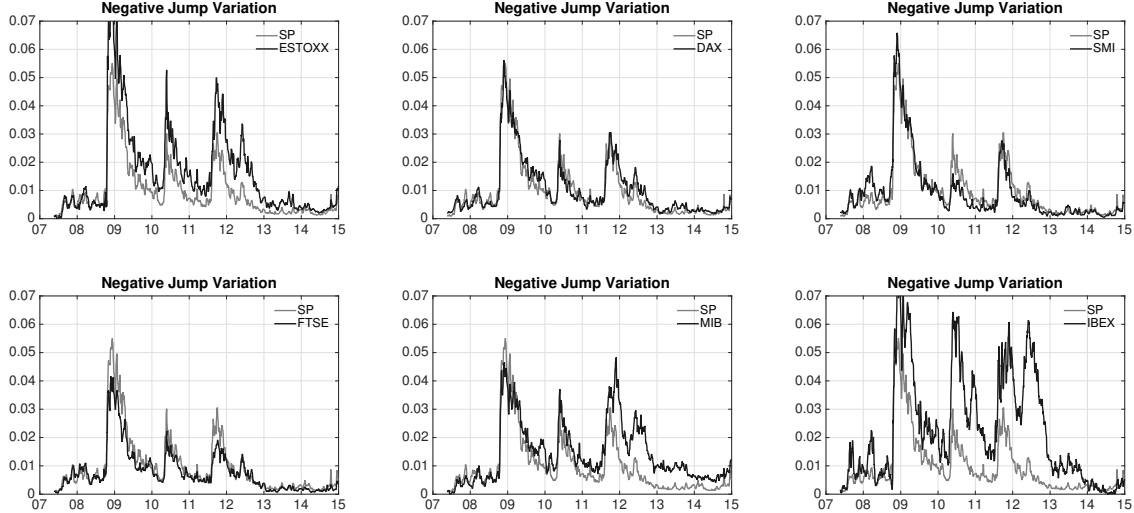


Figure 3: **U Factor Comparison.** For each option-implied negative jump intensity factor, we report the trailing five-day moving average of the implied risk-neutral jump variation $\int_{x<0} x^2 \nu_t^{\mathbb{Q}}(dx) = NJV_t \equiv \frac{2}{\lambda_-^2} U_t$. The series are given in decimals and refer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

spikes in the jump intensity differ substantially from those in volatility, with the sovereign debt crisis inducing a stronger surge in the jump intensity than (diffusive) volatility. We also note that the volatility factors in the U.K. and U.S. evolve in near unison, even if the British jump variation is slightly lower throughout. Similarly, the jump intensity factor for the Swiss index correlates strongly with the S&P 500 factor, although the Swiss factor tends to be above the one in the U.S., both before and during the financial crisis, and then below after the summer of 2009.

We also observe a strong coherence between the S&P 500, ESTOXX and DAX series through the financial crisis and then an elevation in the latter two from the summer of 2009 onwards, with the effect being more pronounced for ESTOXX than DAX, again suggesting a lower exposure of Germany to this crisis. As before, however, the most striking contrast occurs for the Southern European indices. The Italian jump variation rises to a level corresponding to the financial crisis in late 2011, while the Spanish is exceptionally highly elevated during the entire sovereign debt crisis. For these two countries, the jump intensities convey a very different impression of the severity of the crisis. This is not surprising given the widespread speculation at the time that either country might abandon the euro currency. In summary, our decomposition of the primary risk factors documents a substantial increase in return volatility for the Southern European indices along with a further amplification of the negative jump risk, especially for Spain.

4.2 Option Factors as Predictors of Future Risk and Risk Premiums

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

A common concern with predictive regressions is the potential look-ahead bias embedded in the regressors. Several features associated with our set-up effectively mitigate this issue. The option factors are extracted daily, using only the option prices and the high-frequency volatility measure for that given day, along with the parameter estimates, obtained strictly from the varying shape of the option surface. As such, the index returns do not play a direct role in our estimation or factor extraction procedures.³

4.2.1 Predicting Equity Risk

What do the option-implied factors tell us about the risk characteristics of the underlying equity-index? To explore this issue, we regress the future realized return variation, $RV_{t,t+h}$, on the option-implied state variables. The $RV_{t,t+h}$ measure is constructed from intraday returns on the equity-index futures augmented with the squared overnight returns. It provides a good proxy for the risks associated with exposure to the market index, see, e.g., Andersen et al. (2003).

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

We run the predictive regression on a weekly basis, forecasting from 1 to 28 weeks, or roughly 6

³In results, available upon request, we document that the option-implied factors are very similar, and provide qualitatively identical forecast performance, whether the parameters are estimated across the full sample or from an initial year of data only. In the latter case, the parameter vector is fixed at the point estimate based on the early parts of the sample and then used for extracting factors and pricing options over the remaining (out-of-sample) period.

months, into the future. Due to the short sample and varying liquidity, the results can be sensitive to outliers. Extreme observations may be genuine, but may also arise due to measurement errors stemming from periodic illiquidity in the option markets, large option bid-ask spreads during turbulent events, data errors, non-synchronous observations, large standard errors for the nonparametric volatility estimators on days with elevated volatility, and potential model misspecification on days with unusual market stress. Such errors may induce poor identification of the factor realizations. Hence, we winsorize all explanatory variables in the predictive regressions at the 98 percent level, limiting the influence of the 1% extreme negative and positive observations. Importantly, we do not modify the dependent variables, so the multi-horizon excess returns and variation measures appearing as left-hand-side variables incorporate all extreme volatility and return realizations.

The regressions take the form,

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

We reiterate that the risk-neutral expected negative jump variation, NJV_t^\perp , is proportional to the component of U_t orthogonal to the spot variance, as implied by equation (5).

The left panels of Figure 4 reveal, for all our equity indices, that the part of the left jump intensity factor orthogonal to spot volatility has no explanatory power for the ex-post realized return variation. Instead, all predictor power is concentrated in the implied spot variance which, of course, is well known to be a powerful predictor of short-term return volatility. The right panels show that the explained variation is very high and qualitatively similar across all indices. The absence of any auxiliary predictive power in the jump factor is striking. It signifies a stark disconnect between the residual movements in the left side of the implied volatility surface (the part not spanned by volatility) and the riskiness, or variation, of the future returns. The inference in Figure 4, and throughout the paper, is based on Newey-West standard errors with lag length $1.3\sqrt{T}$, as recommended by Lazarus et al. (2018). In Appendix A.2, we provide inference based on two alternative ways to compute the long-run variance of the estimator, following procedures proposed by Lazarus et al. (2018) and Wei and Wright (2011). In either case, the findings are qualitatively identical, even if the standard errors increase slightly, rendering the overall evidence marginally less significant.

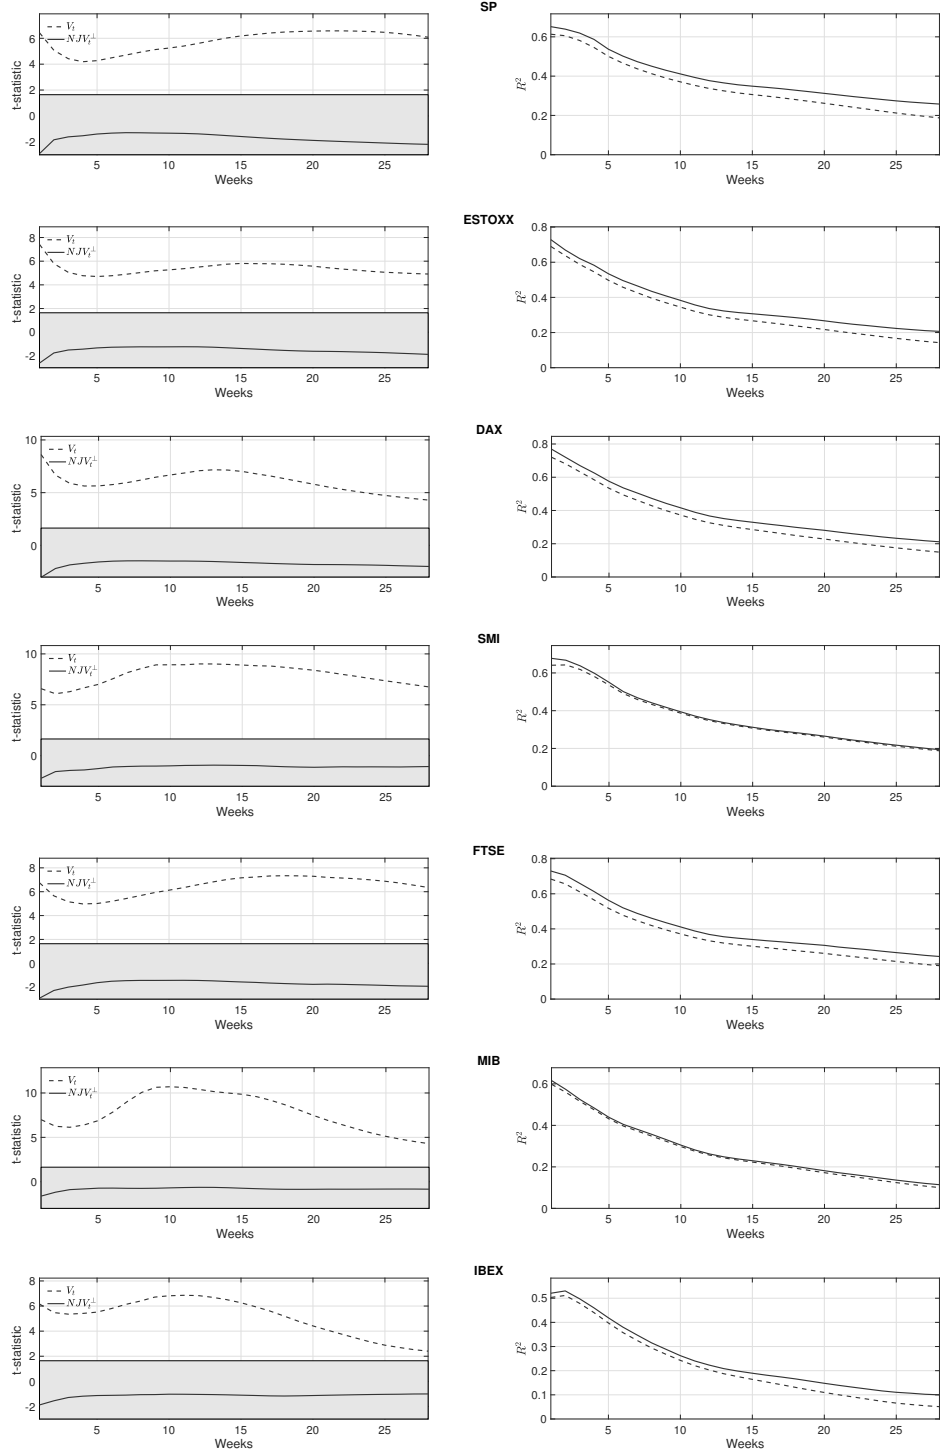


Figure 4: **Predictive Regressions for Return Variation.** Left Panel: t-statistics for the regression slopes according to the Newey-West estimator of the long-run variance with number of lags equal to $1.3\sqrt{T}$; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

4.2.2 Predicting Equity Excess Returns

An important implication of the findings in Section 4.2.1 is that the orthogonal component of the negative jump variation, NJV^\perp , will not be recognized as a factor driving any facet of the risk dynamics in standard time series models estimated from the underlying asset prices. Hence, traditional approaches, identifying the risk factors exclusively from the asset return dynamics, will fail to recognize NJV^\perp as an option pricing factor. In fact, given such findings, one may conjecture that this tail factor is a purely idiosyncratic feature of OTM put option pricing, unrelated to both risk and risk pricing across broader asset markets. It may represent market segmentation arising from frictions and clientele effects in the derivatives markets.

The above hypothesis would imply that the tail factor should possess no auxiliary explanatory power for the pricing of equity risk. We explore this conjecture by running predictive return regressions analogous to those for the return risk in Section 4.2.1. Once again, we rely on mild winsorizing of the regressors, and we continue to ascribe all explanatory power, stemming from joint variation in the state variables, to the traditional spot variance factor.

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

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

Figure 5 conveys the results from the predictive regression (8). There is evidence of a substantial degree of return predictability at the 4-6 month horizon. In addition, we observe a complete role reversal relative to the forecast for the return variation: the significant explanatory variable for future returns is the left tail factor, while the volatility factor is largely irrelevant.

For example, the left panel in the first and second row of Figure 5 plot the t-statistics for the S&P 500 and ESTOXX regression slopes in equation (8), while the right panels display the corresponding R^2 statistics. For both indices the predictive power is low at high frequencies, but rises steadily with the horizon until about five months.⁴ Within the four-month mark, the R^2 surpasses 10%, and it exceeds 15% after six months. The increasing forecast power for longer return horizons is consistent with the hypothesis of a time-varying and persistent equity risk premium, combined with a second mildly persistent return component, that is correlated with the regressors,

⁴This qualitative pattern is familiar, as it is observed for other return predictor variables as well, including the dividend-price ratio and moving averages of interest rates. Yet, there is an important distinction, as these predictors are much more persistent, and they attain significance only at much longer multi-year horizons.

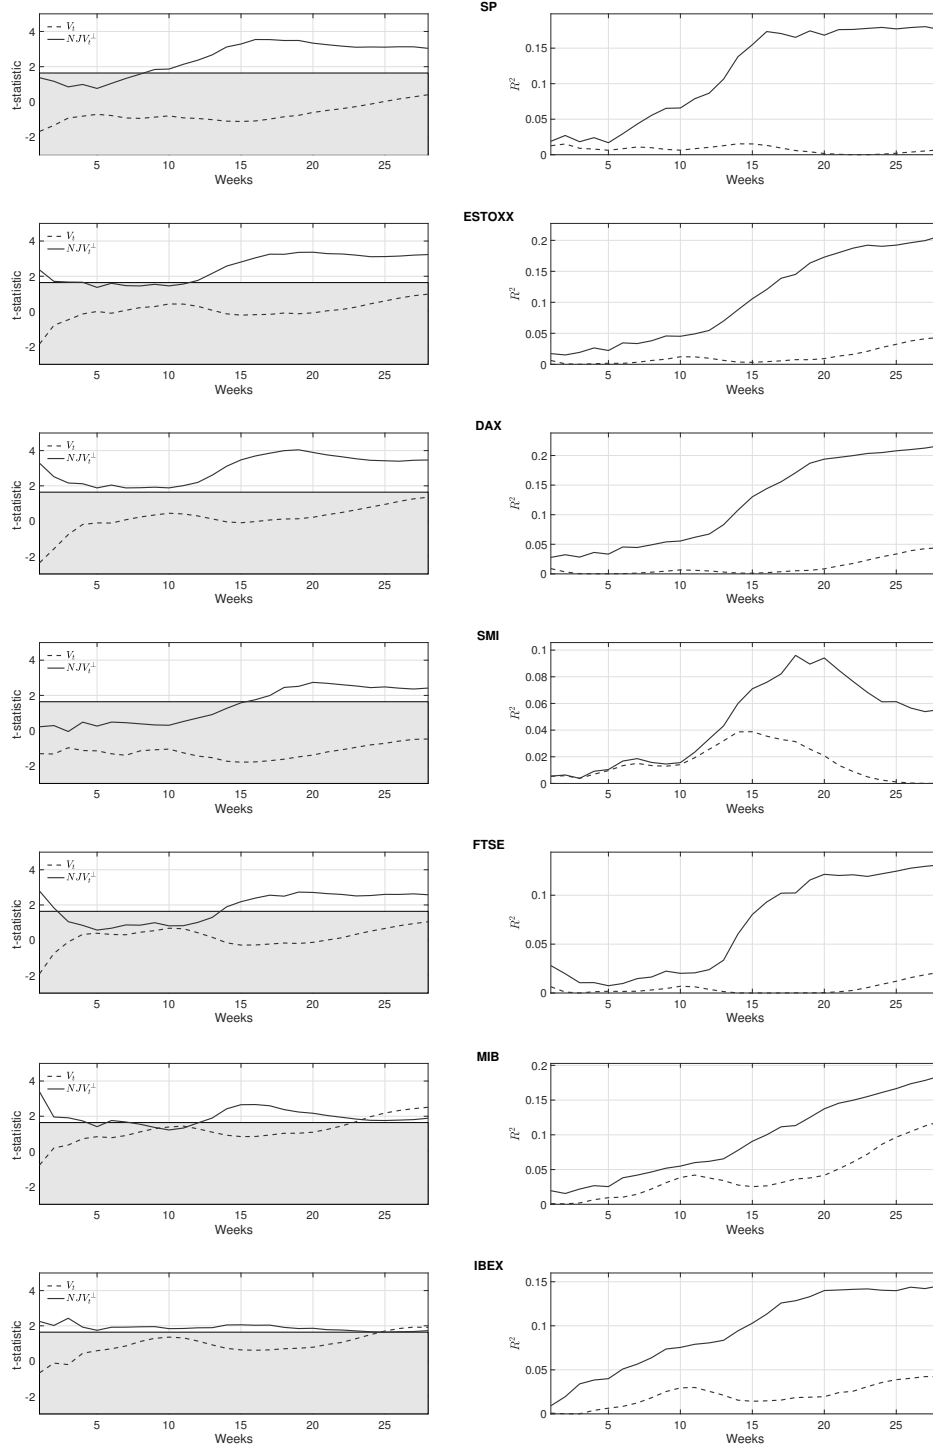


Figure 5: **Predictive Regressions for Excess Returns.** Left Panel: t-statistics for the regression slopes according to the Newey-West estimator of the long-run variance with number of lags equal to $1.3\sqrt{T}$; Right Panel: Regression R^2 , where the full drawn line depicts the total degree of explained variation and the dashed line represents the part explained by the spot variance alone.

see, e.g., the discussion in Stambaugh (1999) and Sizova (2016). The interpretation is that, at the weekly horizon, the largely unpredictable and noisy short-term component dominates, whereas the predictable return component emerges over longer holding periods.

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

Turning to the remaining indices in Figure 5, we observe qualitatively similar features across the board. The only noteworthy differences appear for the two Southern European indices MIB and IBEX, which display a slightly lower degree of significance for the residual left tail factor at the (one-sided) 5% level and some explanatory power for the volatility factor at the longest horizons. Since the evidence for predictability beyond 5-6 months must be viewed somewhat skeptically given the short sample, this is not a surprising finding for two countries subject to extreme dislocations in the market during both the financial and subsequent sovereign debt crises.

Moreover, we note that the degree of explained variation consistently falls between 12-20% at the longer horizons, excluding the Swiss index. For Switzerland, the sudden implementation of an exchange rate cap of the Swiss franc versus the euro on September 6th, 2011, represented a major shock to the equity market, rendering the explanatory power lower than for the other indices, even if the results are not qualitatively different. Specifically, the sharp depreciation of the franc at the introduction of the cap was accompanied by a large positive jump in the (franc denominated) index. Since this intervention was unprecedented, and certainly unexpected, the sharp appreciation of the local index was not reflected a priori in the option surface.⁵

We conclude that the variance factor, effectively, is bereft of explanatory power in these regressions. In contrast, the orthogonalized tail factor provides robust predictive power for the excess returns. This finding is remarkable given the huge discrepancy in the realized index returns over the sample, and the diverse exposures they exhibit vis-a-vis the European debt crises. We also note

⁵Furthermore, since the appreciation of the index was smaller than the simultaneous devaluation of the Swiss franc, the equity market performance, per se, is a poor guide to the return performance viewed from a global perspective. These concerns are even more relevant for the discussion of return predictability associated with the variance risk premium in the following section. The dismantling of the exchange rate cap in 2015 also took the markets by surprise, but it occurred after the end of our sample period.

that the insignificant volatility factor is in line with an extensive time series literature, which has failed to generate consistent evidence that the equity-index return volatility predicts future equity returns, see, e.g., French et al. (1987) and Glosten et al. (1993) for early references. Importantly, our evidence suggests that the option surface does embed critical information for future market returns, but it is contained within factors that are unspanned by volatility.

5 Option Factors and The Pricing of Risk

This section explores the pricing of the different sources of risk and their connection to the extracted option factors more directly.

5.1 The Variance Risk Premium

We start with the variance risk premium, which has been the subject of much interest in the recent literature. The risk-neutral expectation of the quadratic return variation is readily obtained nonparametrically using a portfolio of close-to-maturity options, generating values closely matching the VIX index at the corresponding maturity. Andersen et al. (2015) document that the VIX index provides a close approximation to the risk-neutral expectation of the future return variation, even in the presence of jumps in the return generating process. We follow the procedure of Carr and Wu (2009) for computing the risk-neutral expected return variation for the 30-day horizon, but note that we obtain qualitatively identical results by using our estimated (parametric) risk-neutral model for the return dynamics and the extracted option factors to construct the VIX measure. For the expected objective or statistical return variation, we rely on standard forecasting techniques. The relevant procedures are detailed in Appendix A.3 and A.4.

The extant literature has consistently found large negative variance risk premiums for equity indices. In our sample, the average estimate ranges from -1.6% to -3.9% in annualized volatility terms, where the underlying realized volatility measure, used to compute VRP, corresponds the one reported in Table ?? augmented with the overnight (close-to-open) squared returns to ensure compatibility with the option-implied volatilities. Given the relatively short time span, these estimates are somewhat noisy, but they are entirely in line with prior evidence. Moreover, recent studies document, across international equity indices, that the variance risk premium has predictive power

for the future equity excess returns; e.g., Bollerslev et al. (2009) and Bollerslev et al. (2014).⁶ In our setting, the state vector should determine the risk-neutral expected return variation and simultaneously, as noted above, it provides a good forecast for the future expected return variation (under the statistical, or \mathbb{P} , measure). Consequently, an affine mapping links the state vector to the variance risk premium. Hence, the fact that we obtain significant predictive power for the future equity excess returns from the bivariate regression involving the two (orthogonalized) factors, V and NJV^\perp , for horizons between two to six months in Figure 5 is consistent with prior results.

Nonetheless, our findings raise critical questions regarding the association between the variance and equity risk premiums. In particular, as detailed in Section 4.2.2, V has essentially no explanatory power for the equity risk premium. Instead, the predictive power resides squarely with the pure (orthogonal) negative jump factor, NJV^\perp . This suggests that only certain components of the variance risk premium have explanatory power for the future equity risk premium. This fact has important implications for our understanding of the pricing of equity risk. Hence, below, we explore the impact of the left tail factor, NJV^\perp , in determining the variance and left jump risk premiums, and ultimately in providing information about the dynamics of the equity risk premium.

5.2 The Role of the Left Tail Option Factor

For clarity, we organize the exposition around a stark composite null hypothesis. Our findings suggest—but do not conclusively document—that the pure tail factor, NJV^\perp , solely reflects the investors attitude towards abrupt downside risk exposure, but carries no information regarding the future expected negative jump risk. Hence, we conjecture that this factor constitutes a pure risk premium with no direct relation to actual future equity risks. Moreover, we stipulate that this tail factor provides a good proxy for the overall negative jump risk premium, and that it is the only element of the variance risk premium with significant predictive power for the equity risk premium.

Exploiting the decomposition, $NJV = NJV^\parallel + NJV^\perp$, where NJV^\parallel denotes the part of NJV spanned (linearly) by V , we obtain the following decomposition of the variance risk premium,

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

⁶This issue also relates to the broader literature on predictability for international equity indices, e.g., Harvey (1991), Bekaert and Hodrick (1992), Campbell and Hamao (1992), Ferson and Harvey (1993) and Hjalmarsson (2010). In addition, see Bakshi et al. (2011) for the predictive ability of other option-based volatility measures.

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

Our null hypothesis imposes restrictions on the interaction among the terms in equation (9). In particular, **(A)** if NJV^\perp is a pure jump risk premium component, it should be uncorrelated with the objective expectation terms in the second line, $E_t^{\mathbb{P}}[NJV_{t,t+h}]$ and $E_t^{\mathbb{P}}[CV_{t,t+h} + PJV_{t,t+h}]$. In addition, **(B)** if the variation in NJV^\perp provides a good approximation to the overall variation in the negative jump risk premium, the second term must display a substantially lower degree of variation than NJV^\perp . Finally, **(C)** NJV^\perp should be uncorrelated with the third term, encompassing the risk compensation for continuous and positive jump variation.

Our null hypothesis also induces a rank-ordering on the predictability of equity returns associated with alternative predictors. Specifically, in the standard predictive regression setting,

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

where Y_t signifies the predictor, we obtain the following empirical implications. First, the best forecast variable across all option factors and risk premiums in equation (10), in terms of the degree of explained variation, R^2 , and significance of the slope coefficient, $d_{v,h}$, is the true source of predictability, presumed to be the pure jump factor, **(1)** $Y_t = NJV_t^\perp$. The second best is the negative jump risk premium, **(2)** $Y_t = NJRP_{t,t+h}$. It comprises the sum of the first two terms in the bottom line of equation (9). Since the second component, per assumption, has no predictive power for the equity premium, it merely adds noise, generating an errors-in-variable problem, thus lowering the signal-to-noise ratio and weakening forecast performance. Third, adding the last term from the second line of equation (9), we obtain the full variance risk premium, i.e., **(3)** $Y_t = VRP_{t,t+h}$. This further aggravates the errors-in-variable problem, because the added term displays a great deal of variation over the sample but, under the null hypothesis, possesses no predictive power. Fourth, **(4)** $Y_t = VIX_t$, should perform even worse, as it includes the full risk-neutral return variation, adding NJV^\parallel , CV , and PJV to NJV^\perp , yet it does not include the offsets provided by the corresponding expected (objective) return variation. Therefore, the degree of variation in NJV^\perp is only a small fraction of that for VIX, and the signal-to-noise ratio is now very low. Yet, the VIX index does embed the tail factor, and thus may retain marginal predictive power. Finally, the spot volatility

factor, **(5)** $Y_t = V_t$, does not encompass any element of the NJV^\perp measure, so it must fail to predict the future returns altogether. We already documented this fact empirically in Figure 5.

5.3 The Negative Jump Risk Premium

We now explore the assertions put forth in Section 5.2. First, although we found NJV^\perp to have no predictive power for the total return variation in equation (7), it may still be correlated with individual components of the expected return variation, especially $E_t^{\mathbb{P}}[NJV_{t,t+h}]$.⁷ Consequently, we run an additional set of predictive regressions analogous to equation (7), but with the realized negative jump intensity as regressand. The results, reported in Figure ?? of Appendix A.7, confirm that NJV^\perp has no explanatory power for the future negative jump variation. Moreover, the same holds true for corresponding regressions with the continuous (and positive jump) return variation as regressand. These findings corroborate hypothesis **(A)**, and support the notion that NJV^\perp may constitute a pure risk premium factor.

Second, one may wonder if NJV^\perp is tied closely to the future risk-neutral continuous variation. In theory, the jump intensity determines the expected number of negative jumps, which coincide with those in volatility. However, empirically, this channel is weak. Across all indices, we find that at least 98% of the variation of $CV_{t,t+h}$ is explained by V_t , which is orthogonal to NJV^\perp . In conclusion, NJV^\perp has no meaningful impact on the risk-neutral return variation, except through the negative jump variation. Furthermore, we have established that it is, effectively, unrelated to all expected future objective return variation components. This is fully consistent with hypothesis **(C)**. Hence, we may treat NJV^\perp exclusively as a contributor to the negative jump risk premium.

Third, we gauge the importance of fluctuations in NJV^\perp for the total variation in NJV , where $NJV = NJV^\parallel + NJV^\perp$. A direct assessment follows from the R^2 statistic associated with our original regression of U_t on V_t , where NJV^\perp is proportional to the regression residual. We find that the share of variation stemming from NJV^\perp ranges from 32% for the S&P 500 index to 74% for the Spanish IBEX index, with a mean value close to 50%.⁸ Thus, roughly half of the movement in the option-implied tail factor, NJV , represents shifts in the jump risk premium, not related to the

⁷This may happen if NJV^\perp is correlated with several components but the effects are of opposite sign, so they cancel. Alternatively, the overall correlation may be dominated by the expected continuous variation (CV), so we have low power to detect correlation of NJV^\perp with the remaining smaller contributors to the expected return variation.

⁸The variation shares for the remaining indices are: 59% for ESTOXX; 46% for DAX; 46% for DAX; 39% for SMI; 38% for FTSE; and 53% for MIB.

expectation regarding future return variation. We further note that, in the second line of equation (9), the NJV^{\parallel} term, to a large extent, will be offset by the expected negative jump variation, $E_t^{\mathbb{P}}[NJV_{t,t+h}]$. That is, time-variation in the overall premium, $NJRP$, is dominated by NJV^{\perp} . This conclusion is confirmed by the direct evidence provided in Appendix A.8, and it serves to corroborate hypothesis **(B)**.

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

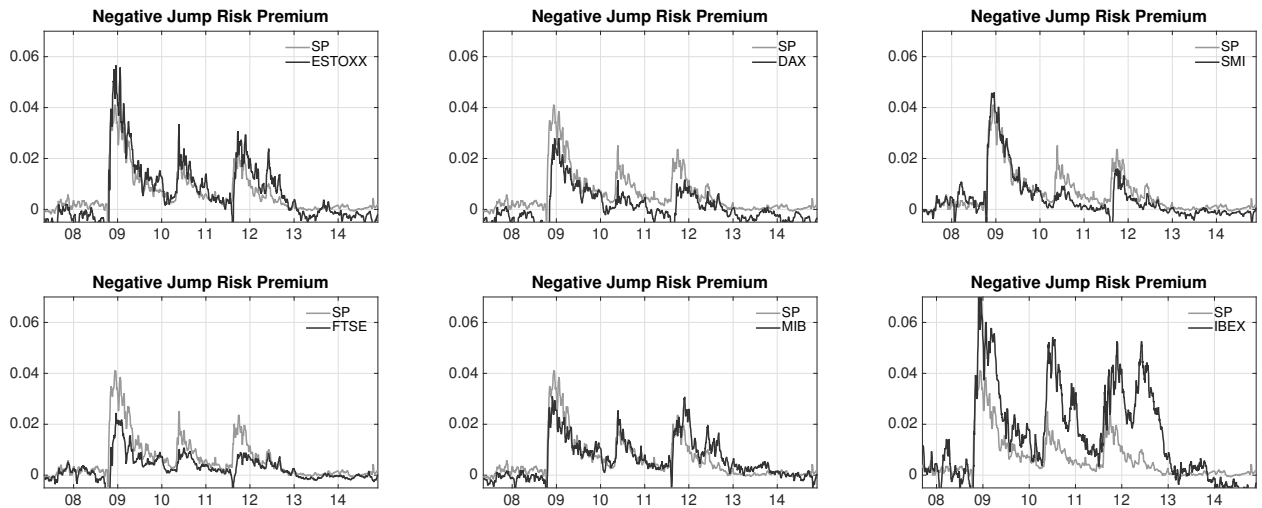


Figure 6: **Negative Jump Risk Premium.** For each index we report the negative jump risk premium as the difference between the risk-neutral and physical expectation of the negative jump variation over thirty calendar days. The series are given in decimals and refer to annualized values.

Figure 6 depicts the time series of estimated negative jump risk premiums for our indices. At any point in time, the premium reflects the extracted jump and volatility factors, the estimates for the risk-neutral return dynamics, and the estimation procedure for the expected negative jump variation. As noted previously, the latter is fairly noisy or imprecise, as conditional jump intensities

and sizes are notoriously difficult to forecast. Thus, there is some uncertainty associated with these series, but the fact that NJV often is an order of magnitude larger than the objective expectation of the negative jump variation renders the pronounced systematic variation in the series reliable and significant.

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

5.4 Risk Premiums as Predictors of Future Excess Returns

In this section, we directly explore the hypothesis that the variance risk premium possesses explanatory power for future equity returns solely because it incorporates the negative jump risk premium. We further relate our findings to the rank-ordering of the alternative predictors, labeled **(1)(5)**, for the equity risk premium outlined in Section 5.2.

In Table 1, we report on predictive regressions that contrast the forecast power of the variance risk premium with that of the negative jump risk premium. The table refers to (constrained or unconstrained) variants of the following regression,

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

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

The results for forecast horizons of 1, 5, and 7 months are summarized in the separate panels of Table 1. At the one-month horizon, the results are largely insignificant and the explained variation

	SP	ESTOXX	DAX	SMI	FTSE	MIB	IBEX						
Panel A: $R_{t,t+1}$													
VRP	-0.42	-1.33	1.07	-0.72	1.51	0.57	-2.27	0.80	0.35	0.12	-1.84	0.71	-1.48
$VRP - NJRP$	-0.78	1.33	0.45	-0.39	2.07	1.74	-4.38	0.44	0.40	0.93	3.24	2.23	-0.36
$NJRP$	0.00	0.01	0.02	0.00	0.01	0.00	0.02	0.02	0.00	0.01	0.00	0.04	0.00
R^2													
Panel B: $R_{t,t+5}$													
VRP	0.49	-1.93	0.85	-2.99	1.81	-0.56	-3.13	0.70	-0.17	0.43	-1.32	0.43	-2.53
$VRP - NJRP$	0.04	2.45	1.17	-1.86	2.71	3.02	-7.07	2.12	1.94	1.81	3.39	2.33	-0.81
$NJRP$	0.01	0.00	0.05	0.03	0.01	0.01	0.11	0.10	0.00	0.01	0.00	0.11	0.09
R^2													
Panel C: $R_{t,t+7}$													
VRP	2.18	-0.88	2.20	-0.77	2.91	0.42	-2.77	1.35	0.20	2.10	0.04	1.30	-1.22
$VRP - NJRP$	1.73	2.93	2.57	-0.41	3.19	3.70	-1.56	2.12	2.71	2.97	1.72	2.27	2.55
$NJRP$	0.06	0.03	0.12	0.12	0.05	0.01	0.16	0.06	0.00	0.12	0.04	0.16	0.14
R^2													

Table 1: Predictive Regressions and Risk Premiums. For each index, the first column reports the t-statistic and the R^2 from the univariate regression of future returns on VRP , while the following three columns report the identical statistics for related regressions of future returns on risk premiums. t-statistics are computed using the Newey-West estimator of the long-run variance with number of lags equal to $1.3\sqrt{T}$. In the second column, the single regressor is the variance risk premium minus the risk premium associated with negative jumps, the third column refers to a bivariate setting featuring both VRP and $NJRP$ as regressors, and in the fourth column the single regressor is the negative jump risk premium. Panel A displays the results for the future one-month, Panel B for future five-month, and Panel C future seven-month returns.

(R^2) is uniformly low, as expected, given the prior evidence of predictability only at longer horizons. Nonetheless, we note that t-statistics for VRP in column one are uniformly lower than those for $NJRP$ in column four. Likewise, in the bivariate regressions, the t-statistic for $NJRP$ is always larger than the one for VRP . The latter feature, indicating superiority of the $NJRP$ relative to VRP , is a robust feature observed across all indices and horizons.

At the five-month horizon, where we establish significant predictive power of our pure tail factor across all indices in Figure 5, the VRP remains insignificant, but the regression coefficient is positive for all indices except the Swiss. Most importantly, $NJRP$ is now significant in the majority of cases, and the reduction in explanatory power, as we move from the bivariate regression to the univariate $NJRP$ regression, is quite limited, except for the Swiss index. Moreover, the R^2 values are moderately high at 7% – 12% for five of the indices. Finally, the evidence is, if anything, weaker for the U.S. than for most of the other indices, implying that the type of results originally suggested by work on the S&P 500 index also applies for the main European markets.

The above tendencies only strengthen at the seven-month horizon. Except for the Swiss index, the R^2 rise considerably, and the majority of the $NJRP$ coefficients are significant. In addition, excluding VRP from the bivariate regression has a negligible impact in terms of explanatory power. This is consistent with the extremely low R^2 statistics for the univariate $VRP - NJRP$ regression, where the negative jump risk premium is stripped from the variance risk premium. This effectively annihilates the predictive power of the VRP measure. Again, to complement the Newey-West standard errors, we provide corresponding evidence based on alternative estimates for the long-run variance in Appendix A.2. The findings verify the relative performance of predictors **(2)** and **(3)**.

In summary, the forecast power for the future equity risk premium resides with the negative jump risk premium, while the variance risk premium has negligible explanatory power, once the contribution from risk pricing of the left tail is netted out. These conclusions mirror our findings regarding the relative predictive power of the tail factor NJV^\perp and volatility factor V in Figure 5, although they appear somewhat less significant and consistent in Table 1. The poor forecast performance for the Swiss index may be explained, in part, by the imposition of a cap on the Swiss franc-euro exchange rate on September 6, 2011, which induced an unexpected large jump in the equity index. This event is highly influential for the regression coefficients governing the forecasts of the future objective return variation. This suggests that the option-implied factors should possess

superior predictive power relative to the risk premiums due to the incremental noise associated with the forecasts for the future return variation. This is consistent with the discrepancy between the results for the SMI index in Figure 5 and Table 1.

We also confirm that our remaining hypotheses, rank-ordering the explanatory power of alternative predictors vis-a-vis the equity premium, hold up. Comparing Table 1 and Figure 5, we find that NJV^\perp performs on par, and in some cases dramatically better, at the 5-month horizon than the $NJRP$, so predictor (1) dominates predictor (2). At one month, the results are similar, but both measures have negligible predictive power. Finally, for the 7-month horizon, where the inference may be a bit less reliable, the two predictors perform essentially on par. In all cases, the VRP produces substantially worse forecasts. Finally, we verify in Appendix A.9 that predictor (4), the VIX index, has trivial explanatory power and only sporadically turns significant at the longest horizons. Although the associated R^2 -statistics, as expected, do exceed those for predictor (5), the V factor, both variables provide ineffective forecasts of the equity risk premium.

Overall, the qualitative results are remarkably robust, in spite of the large discrepancies in cumulative returns highlighted in Figure 1. In other words, the predictive association between option-implied factors and the future realized returns and return variation appears to be operative almost uniformly for the U.S. and major European equity-index and derivatives markets.

6 On Commonality in Risk and Risk Pricing

Even though country-specific features are evident, our results generally suggest a substantial degree of coherence across the markets, as is also evident in Figures 1–3. Inspired by these observations, we now briefly explore the covariation between the risks and risk premiums across our equity markets. Table 2 reports, for two subsamples, the pairwise correlation of the expected negative jump and continuous return variation, with the latter approximated by the total variation net of the negative jump variation. These objective risk measures display a high degree of coherence. The most notable feature is a relatively weak association between the Italian NJV and the other indices in the first subsample and a reduction in the correlation of the Spanish return variation with the remaining indices in 2010-2014 relative to 2007-2009. In sum, the individual indices appear to confront similar objective jump and volatility risk perceptions, with some indications that the Italian and Spanish indices display different risk exposures.

		$NJV_{t,t+30}^{\mathbb{P}}$													
		2007-2009					2010-2014								
		SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
SP500		96		0.91	0.85	0.95	0.90	0.84		0.94	0.88	0.72	0.95	0.79	0.68
ESTOXX				0.98	0.92	0.99	0.86	0.93			0.95	0.80	0.97	0.88	0.84
DAX					0.95	0.98	0.78	0.96				0.88	0.95	0.84	0.73
SMI						0.94	0.66	0.97					0.85	0.61	0.65
FTSE							0.81	0.94						0.81	0.78
MIB															
IBEX															0.72

		$QV_{t,t+30}^{\mathbb{P}} - NJV_{t,t+30}^{\mathbb{P}}$													
		2007-2009					2010-2014								
		SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
SP500		0.98		0.97	0.96	0.99	0.92	0.93		0.96	0.90	0.88	0.97	0.82	0.74
ESTOXX				0.99	0.96	0.99	0.95	0.97			0.95	0.90	0.98	0.93	0.83
DAX					0.95	0.98	0.93	0.96				0.96	0.94	0.91	0.68
SMI						0.97	0.91	0.95					0.92	0.82	0.60
FTSE							0.93	0.96						0.87	0.77
MIB															
IBEX								0.89							0.85

Table 2: NJV and $RV - NJV$ Correlation

		$NJRP_{t,t+30}$													
		2007-2009					2010-2014								
		SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
SP500		0.94	0.73	0.98	0.98	0.77	0.95	0.94		0.91	0.62	0.84	0.83	0.87	0.80
ESTOXX			0.90	0.94	0.94	0.91	0.97	0.97			0.85	0.90	0.95	0.95	0.92
DAX				0.75	0.95	0.95	0.82	0.86				0.84	0.88	0.80	0.75
SMI					0.78	0.92	0.94	0.94					0.88	0.83	0.73
FTSE						0.82	0.86	0.86						0.86	0.90
MIB							0.95	0.95							0.89
IBEX															

		$VRP_{t,t+30} - NJRP_{t,t+30}$													
		2007-2009					2010-2014								
		SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX	SP500	ESTOXX	DAX	SMI	FTSE	MIB	IBEX
SP500		0.94	0.90	0.93	0.93	0.70	0.96	0.90		0.19	0.19	0.24	0.46	0.80	0.82
ESTOXX			0.94	0.98	0.98	0.82	0.97	0.97			0.66	-0.12	0.30	0.29	0.25
DAX				0.92	0.92	0.82	0.94	0.93				0.24	0.44	0.16	0.21
SMI					0.78	0.97	0.99	0.99					0.42	0.21	0.13
FTSE						0.72	0.76	0.76						0.40	0.39
MIB							0.96	0.96							0.80
IBEX								0.96							0.80

Table 3: $NJRP$ and $VRP - NJRP$ Correlation

Table 3 speaks to the degree of commonality in the pricing of risk by reporting pairwise correlations of the estimated premiums associated with the one-month left jump tail and continuous return variation, respectively. The top panel shows that the compensation for negative jump risk is highly correlated across the indices in both sub-periods, with only the IBEX deviating moderately in the second period, reflecting the elevated downside risk pricing in Spain during the sovereign debt crises. In contrast, the correlation of the premium associated with the continuous return variation, proxied by the total variance risk premium minus the contribution of the negative jump risk premium, changes substantially across the two subsamples. Initially, during 2007-2009, the compensation for (diffusive) variance risk is near identical across indices, as indicated by the very high correlations reported in the left side of the lower panel of Table 3. In contrast, many pairwise correlations are almost negligible in the subsequent period. Hence, as the sovereign debt crisis unfolds, the option surfaces display more distinct country-specific features, reflecting much stronger heterogeneity in the pricing of (continuous) variance risk.

Given the divergence in the broader risk pricing over the last subsample, the coherence in the left tail risk pricing is remarkable. It suggests a strong degree of commonality in the attitude towards downside equity risk, generating a strong correlation in equity market performance. In fact, it is evident from Figure 1 that the return correlation is very strong, i.e., the periods of market appreciation and depreciation are highly synchronized. Hence, generally speaking, while some indices are subject to larger and more frequent negative shocks than others, there is a striking similarity in the response to such shocks across markets. In particular, they have roughly analogous consequences for risk pricing, as manifest in their impact on option valuation across the moneyness-maturity spectrum, and consequently also on their extracted factor realizations. Through this channel, they induce a high degree of correlation in signals for the direction of the future index appreciation and the associated return variation.

Our findings suggest that the equity markets are well integrated. Nonetheless, a formal analysis of global pricing of left tail risk must involve an explicit consideration of foreign exchange exposures and risk pricing as well, because this requires an assessment of the risk and risk exposures in a common currency unit. For downside tail exposures, this issue is particularly pertinent, as key currency values are known to fluctuate systematically in response to global economic shocks.

7 Conclusion

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

We further document that the variance risk premium only has predictive power for the future equity returns due to the inclusion of the negative jump risk premium within the measure. We also show that the left tail factor and negative jump risk premium remain highly correlated across indices in the period following the financial crisis and through the European sovereign debt crises. In contrast, the volatility factor and diffusive volatility risk premium display a sharp drop in correlation across indices in the second part of the sample. This suggests a strong degree of commonality in the pricing of equity risk internationally, linked to the relative strength of the risk-neutral left jump intensity, whereas the underlying market risks at times vary markedly.

References

- Almeida, C., K. Ardison, R. Garcia, and J. Vicente (2017). Nonparametric Tail Risk, Stock Returns and the Macroeconomy. *Journal of Financial Econometrics* 15, 333–376.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and Forecasting Realized Volatility. *Econometrica* 71, 579–625.
- Andersen, T. G., O. Bondarenko, and M. T. Gonzalez-Perez (2015). Exploring Return Dynamics via Corridor Implied Volatility. *Review of Financial Studies* 28, 2902–2945.
- Andersen, T. G., N. Fusari, and V. Todorov (2015a). Parametric Inference and Dynamic State Recovery from Option Panels. *Econometrica* 83, 1081–1145.
- Andersen, T. G., N. Fusari, and V. Todorov (2015b). The Risk Premia Embedded in Index Options. *Journal of Financial Economics* 117, 558–584.

- Andersen, T. G., N. Fusari, and V. Todorov (2017). Pricing Short-Term Market Risk: Evidence from Weekly Options. *Journal of Finance* 72, 1335–1386.
- Atilgan, Y., T. Bali, and K. Demirtas (2013). Implied Volatility Spreads and Expected Market Returns. *Journal of Business and Economic Statistics* 33(1), 87–101.
- Bakshi, G., G. Panayotov, and G. Skoulakis (2011). Improving the Predictability of Real Economic Activity and Asset Returns with Forward Variances Inferred from Option Portfolios. *Journal of Financial Economics* 100, 475–495.
- Bali, T., K. Demirtas, and H. Levy (2009). Is There an Intertemporal Relation between Downside Risk and Expected Returns? *Journal of Quantitative and Financial Analysis* 44, 883–909.
- Bali, T. and A. Hovakimian (2009). Volatility Spreads and Expected Stock Returns. *Management Science* 55(11), 1797–1812.
- Bali, T. and S. Murray (2013). Does Risk-Neutral Skewness Predict the Cross Section of Equity Option Portfolio Returns? *Journal of Quantitative and Financial Analysis* 48(4), 1145–1171.
- Bates, D. S. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies* 9, 69–107.
- Bates, D. S. (2000). Post-'87 Crash Fears in S&P 500 Future Options. *Journal of Econometrics* 94, 181–238.
- Bates, D. S. (2003). Empirical Option Pricing: A Retrospection. *Journal of Econometrics* 116, 387–404.
- Bekaert, G. and R. Hodrick (1992). Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets. *Journal of Finance* 47, 467–509.
- Bollerslev, T., J. Marrone, L. Xu, and H. Zhou (2014). Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence. *Journal of Financial and Quantitative Analysis* 49, 633–661.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T. and V. Todorov (2011). Tails, Fears and Risk Premia. *Journal of Finance* 66, 2165–2211.
- Bollerslev, T., V. Todorov, and L. Xu (2015). Tail Risk Premia and Return Predictability. *Journal of Financial Economics* 118, 113–134.

- Broadie, M., M. Chernov, and M. Johannes (2007). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance* 62, 1453–1490.
- Campbell, J. and Y. Hamao (1992). Predicting Stock Returns in the United States and Japan: A Study of Long-Term Capital Market Integration. *Journal of Financial Markets* 47, 43–69.
- Carr, P. and L. Wu (2009). Variance Risk Premia. *Review of Financial Studies* 22, 1311–1341.
- Christoffersen, P., K. Jacobs, and K. Mimouni (2010). Models for S&P 500 Volatility Dynamics: Evidence from Realized Volatility, Daily Returns, and Option Returns. *Review of Financial Studies* 23, 3141–3189.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai (2012). Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options. *Journal of Financial Economics* 106, 447–472.
- Conrad, J., R. Dittmar, and E. Ghysels (2013). Ex ante Skewness and Expected Stock Returns. *Journal of Finance* 68(1), 85–124.
- Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics* 7, 174–196.
- Du, J. and N. Kapadia (2012). Tail and Volatility Indices from Option Prices. Working paper.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.
- Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Ferson, W. and C. Harvey (1993). The Risk and Predictability of International Equity Returns. *Review of Financial Studies* 6, 527–566.
- French, K., W. Schwert, and R. Stambaugh (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics* 19, 3–29.
- Gao, G., X. Lu, and Z. Song (2018). *Management Science*, forthcoming.
- Glosten, L., R. Jagannathan, and D. Runkle (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance* 48, 1779–1801.
- Gruber, P., C. Tebaldi, and F. Trojani (2015). The Price of the Smile and Variance Risk Premia. Working paper, University of Lugano and University of Geneva.

- Harvey, C. (1991). The World Price of Covariance Risk. *Journal of Finance* 46, 111–157.
- Hjalmarsson, E. (2010). Predicting Global Stock Returns. *Journal of Financial and Quantitative Analysis* 45, 49–80.
- Hodrick, R. (1992). Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement. *The Review of Financial Studies* 5(3), 357–386.
- Jones, C. (2003). The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets. *Journal of Econometrics* 116, 181–224.
- Kelly, B. and H. Jiang (2014). Tail Risk and Asset Prices. *Review of Financial Studies* 27, 2841–2871.
- Kou, S. (2002). A Jump- Diffusion Model for Option Pricing. *Management Science* 48, 1086–1101.
- Kozhan, R., A. Neuberger, and P. Schneider (2013). The Skew Risk Premium in the Equity Index Market. *Review of Financial Studies* 26(9), 2174–2203.
- Lazarus, E., D. Lewis, J. Stock, and M. Watson (2018). HAR Inference: Recommendations for Practice. *working paper*.
- Li, J. and G. Zinna (2018). The Variance Risk Premium: Components, Term Structures, and Stock Return Predictability. *Journal of Business Economics and Statistics* 36, 411–425.
- Londono, J. (2016). The Variance Risk Premiums around the World. Working paper.
- Londono, J. and H. Zhou (2017). Variance Risk Premiums and the Forward Premium Puzzle. *Journal of Financial Economics* 124, 415–440.
- Merton, R. (1976). Option Pricing when Underlying Asset Returns are Discontinuous. *Journal of Financial Economics* 3, 125–144.
- Mueller, P., P. Sabtchevsky, A. Vedolin, and P. Whelan (2017). Variance Risk Premia on Stocks and Bonds. Working paper.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Santa-Clara, P. and S. Yan (2010). Crashes, Volatility and the Equity Premium: Lessons from S&P 500 Options. *Review of Economics and Statistics* 92, 435–451.
- Sizova, N. (2016). Do Long-Run Relations Convey Extra Information? Working paper, Rice University.

Stambaugh, R. (1999). Predictive Regressions. *Journal of Financial Economics* 54, 783–820.

Wang, H., H. Zhou, and Y. Zhou (2013). Credit Default Swap Spreads and Variance Risk Premia. *Journal of Banking and Finance* 37, 3733–3746.

Wei, M. and J. Wright (2011). Reverse Regression and Long Horizon Forecasting. *Journal of Applied Econometrics* 28(3), 353–371.