High-Frequency Tail Risk Premium and Stock Return Predictability *

Caio Almeida[†] Kym Ardison[‡] Gustavo Freire[§] René Garcia[¶] Piotr Orłowski[‖]

May 16, 2023

Abstract

We propose a novel measure of the market return tail risk premium based on minimum-distance state price densities recovered from high-frequency data. The tail risk premium extracted from intra-day S&P 500 returns predicts the market equity and variance risk premiums and expected excess returns on a cross section of characteristics-sorted portfolios. Additionally, we describe the differential role of the quantity of tail risk, and of the tail premium, in shaping the future distribution of index returns. Our results are robust to controlling for established measures of variance and tail risk, and of risk premiums, in the predictive models.

Keywords: Tail Risk, Risk-Neutral Measure, Expected Shortfall, Intra-day Market Returns, Return Predictability

JEL Code: G12, G13, G17.

*We would like to thank Rodrigo Hizmeri, seminar participants at the Kellogg School of Management and conference participants at the 2017 SoFiE Conference in New York, the 2017 Vienna-Copenhagen Conference on Financial Econometrics, and the 2018 IAAE Meeting in Montreal for useful comments and suggestions. The second author acknowledges financial support from ANBIMA and FAPERJ. The fourth author thanks the NSERC, the SSHRC and the FQRSC research grant agencies for their financial support. He is a TSE associate faculty and a research Fellow of CIRANO and CIREQ.

†Corresponding Author. Email: calmeida@princeton.edu, Department of Economics and Bendheim Center for Finance, Princeton University, phone: 609-424-4203.

‡Email: kymmarcel@gmail.com, SPX Capital, Rio de Janeiro, Brazil.

§Email: freire@ese.eur.nl, Erasmus University Rotterdam and Tinbergen Institute.

¶Email: rene.garcia@umontreal.ca, Université de Montréal and Toulouse School of Economics.

‖Email: piotr.orlowski@hec.ca, HEC Montréal, CDI Associate Fellow.

1 Introduction

Starting with Bollerslev et al. (2009), extant empirical evidence supports that the variance risk premium (VRP) helps predict future aggregate market returns.[1] This predictability has been rationalized under different theoretical frameworks.[2] Focusing on the properties of the VRP in a model-free manner, Bollerslev and Todorov (2011) and Bollerslev et al. (2015) demonstrate how this premium reflects compensation for two different types of risk: diffusive and jump risk. They show that a large part of the VRP and its predictive power for the equity premium comes from the compensation demanded by investors for extreme negative events. Such asymmetric importance of losses relative to gains has strong theoretical foundations, as in the case of loss aversion (Kahneman and Tversky, 1979) and disappointment aversion (Gul, 1991; Routledge and Zin, 2010) preferences. These theories posit that agents are specially averse to downside losses, such that greater compensation (in the form of higher expected returns) is demanded for assets with high downside risk.[3]

In this paper, we introduce a new measure of the compensation that investors demand for bearing systematic downside risk. We use it to shed light on the predictability of aggregate and cross-sectional risk premiums that is due to aversion to downside risk in a high-frequency environment. The motivation of our analysis is threefold. First, theory indicates that the VRP should predict market returns, where most of the VRP arises directly from compensation for downside risk. Second, theory also suggests that investors will demand higher

[1]The VRP is defined as the difference between the conditional expected variance of market returns over a given horizon under the physical and risk-neutral measures. It captures the compensation demanded by investors for bearing variance risk.

[2]Bollerslev et al. (2009) and Drechsler and Yaron (2011) extend the long-run risk model of Bansal and Yaron (2004) to make the equity premium a function of time-varying volatility-of-volatility and jump-intensity, respectively. In both cases, the VRP effectively isolates the relevant latent factor (vol-of-vol or jump-intensity), justifying its predictive power for market returns. This is because the representative agent's aversion to a shock in the latent factor makes the risk-neutral variance higher than the physical one, i.e., the VRP increases (in absolute value) with the factor. Alternatively, Bonomo et al. (2015) show that the predictability afforded by the VRP can be generated by incorporating generalized disappointment aversion preferences (Routledge and Zin, 2010) in the long-run risk model. See also Bekaert and Engstrom (2017), who extend the habit formation preferences of Campbell and Cochrane (1999) to show that the VRP can be interpreted as a proxy for aggregate risk aversion.

[3]This prediction has been empirically confirmed for the cross-section of stock returns by Ang et al. (2006) and Farago and Tédongap (2018).

returns when their perception of downside risk is high. Third, while the predictability literature focuses on monthly to longer horizons, high-frequency evidence can provide new stylized facts for economic theories to account for, and to be judged against.[4]

We compute the daily tail risk premium, denoted as $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, as the difference between the expected shortfall calculated under the risk-neutral ($ES^{\mathbb{Q}}$) and physical ($ES^{\mathbb{P}}$) probability measures estimated from intra-day S&P 500 return data. The risk-neutralization is based on a nonparametric adjustment of the raw market returns. Motivated by Ait-Sahalia and Lo (1998, 2000), the risk adjustment puts a higher probability weight on extreme negative returns, reflecting investors' compensation for "bad" states of the world. From an economic viewpoint, $ES^{\mathbb{Q}}$ ($ES^{\mathbb{P}}$) is the expectation under the risk-neutral (physical) measure of the payoff of a hypothetical out-of-the-money (OTM) put option on the market, which is naturally sensitive to negative jumps in returns. In particular, $ES^{\mathbb{Q}} - ES^{\mathbb{P}}$ can be expressed as the expected value of the difference between the put price and its payoff, which allows us to interpret $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ as the expected gain of selling the put. Therefore, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ will be high when investors highly value protection against market downside risk (i.e., when they are willing to pay a high premium relative to the expected put payoff).

We investigate the predictive power of $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ for risk premiums with a comprehensive set of empirical tests. We focus our predictive exercises on very short horizons: daily, weekly and monthly. This directly complements the longer-horizon predictability results documented in the literature. In the predictive regressions, we control for a number of established variance and tail risk measures. Among these, three are particularly important: the physical expected shortfall ($ES^{\mathbb{P}}$), the left tail variance (LTV) of Bollerslev et al. (2015) and the VRP of Bollerslev et al. (2009). Including $ES^{\mathbb{P}}$, which is a measure of realized downside risk, allows us to assess what better predicts risk premiums: the tail risk premium or the quantity of tail risk itself. LTV, which estimates jump risk from out-of-the-money options, is a natural benchmark to assess if options contain information beyond that provided by our measure.

[4]For instance, Bonomo et al. (2015) develop a model to reproduce moments and predictability patterns of risk and return across different frequencies.

The VRP is included given its important role as a predictor of the equity premium (Bollerslev et al., 2009).[5] We complement these measures with the realized variance (RV) of market returns computed from intra-day data and other measures of realized return variation.

We start with a predictive analysis of excess market returns. We find that our measure $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ has strong predictive power for 1-day ahead market returns, with a positive coefficient that is highly statistically significant and a non-negligible R^2.[6] This provides new high-frequency evidence that investors require a higher compensation to hold the market when aversion to downside risk increases. This predictive power is robust to controlling for all the measures we consider, across different regression specifications. In particular, the only other predictor that appears as significant is the VRP. We show that the predictability afforded by $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ also holds out-of-sample across different starting dates and re-estimation frequencies. This indicates that the predictive relation between the tail risk premium and the equity premium is stable and persists under different economic conditions.

We further investigate whether $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is able to predict the variance risk premium. Since the VRP reflects in large part compensation for extreme negative events, we should expect that the tail risk premium is informative about the future VRP. This is indeed what we find in the predictive regressions, where $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is a highly significant predictor across different regression specifications and all horizons. In particular, this predictability remains after including the lagged VRP, which is also highly significant. It is worth noting that the LTV has no predictive power for the equity and variance risk premiums. This can be rationalized by the fact that it is computed from options with maturity between 6 and 31 trading days, which reflect market expectations over relatively long horizons. In contrast, our tail risk premium measure contains information completely conditional on day t.[7]

We also analyze the predictive relation between the tail risk premium and 1-day ahead

[5]We compute a daily version of the VRP that is analogous to the monthly VRP constructed by (Bollerslev et al., 2009).

[6]For the 1-week and 1-month ahead horizons, no predictor is significant.

[7]In Appendix A, we show that, similarly to LTV, our tail risk premium measure captures time-variation in the risk-neutral and physical tail shape parameters of the market return distribution. However, the two measures naturally differ in how they are estimated and what data is used.

excess returns on portfolios comprised of stocks sorted on various characteristics: size, book-to-market, profitability, investment, momentum, reversal and industry. These portfolios reflect compensation for risks beyond those related to the market portfolio. Given the important role of downside risk factors in explaining the cross-section of returns (Ang et al., 2006; Farago and Tédongap, 2018), $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ can potentially be useful for predicting cross-sectional risk premiums at high-frequency. Our results show that the tail premium measure is able to predict the returns on most of the portfolios with high t-statistics and economically significant R^2s. The same is not true for the other variables we consider. In fact, adding each of them at a time to the predictive regressions does not affect the predictive power of $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$. This reinforces the role of aversion to downside risk as captured by our measure as a fundamental determinant of risk premiums at short horizons.

Our analysis indicates that it is the tail risk premium ($\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$) rather than the level of tail risk ($ES^{\mathbb{P}}$) that contains relevant predictive information about future market returns.[8] In order to shed light on the differential role of these variables for explaining the equity premium, we investigate their predictive power for the whole distribution of next-day market returns by estimating the quantile regression model of Koenker and Gilbert (1978). We find that an increase in risk ($ES^{\mathbb{P}}$) leads to a larger probability of observing both extreme negative and positive market returns, whereas an increase in the aversion to downside risk ($\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$) shifts the quantiles around the median and the whole right tail towards more positive values. That is, a positive shock in expected shortfall means a more volatile market, such that it is usually followed by either a large decrease or increase of the S&P 500 index. These extreme effects cancel out when predicting directly the market returns. In contrast, a positive shock in the tail risk premium signals that investors are more averse to extreme negative outcomes, requiring a higher compensation to hold the market. This is reflected in the positive effect of $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ on essentially all quantiles of the market return distribution. Such unambiguous effect

[8]This result is largely in line with Bollerslev et al. (2009), Bollerslev and Todorov (2011), Bollerslev et al. (2015) and Andersen et al. (2017), among others, who study the risk-return trade-off concluding that it is the variance risk premium rather than the quantity of risk that is informative of future risk premiums.

translates to the significant positive relation between $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and future market returns.[9]

The rest of the paper is organized as follows. Section 1.1 offers a brief review of the related literature. Section 2 describes how we define and estimate our tail risk premium measure based on high-frequency market returns. Section 3 presents our empirical predictability results for risk premiums and quantiles of the S&P 500 excess return distribution. Section 4 contains several robustness tests considering different specifications of our tail premium measure. Section 5 concludes the paper.

1.1 Related Literature

Our paper is related to the rich and growing literature on the estimation of tail risk and systematic risk measures and their use to predict the equity and variance risk premiums. This includes, among others, Bali et al. (2009), Allen et al. (2012), Siriwardane (2013), Kelly and Jiang (2014), Adrian and Brunnermeier (2016) and Brownlees and Engle (2017). In particular, Bollerslev et al. (2015) decompose the variance risk premium and examine the importance of the diffusive and jump components for return predictability. Andersen et al. (2017) estimate the variation in the left tail of the return distribution from short-maturity options with significant predictive power for future short-term returns. Andersen et al. (2020) extend this evidence to international equity markets. Vilkov and Xiao (2013), Ghysels and Wang (2014) and Huggenberger et al. (2018) use daily index options to model forward looking tail risk based on Value-at-Risk and expected shortfall measures. In contrast to these papers, we are the first to provide a methodology to compute the tail risk premium in high-frequency environments that is applicable to virtually any set of returns.

To estimate the risk-neutral leg of our tail risk premium measure, we build on Almeida et al. (2017), who compute a nonparametric risk-neutral expected shortfall based on a cross-section of daily portfolio or security returns.[10] In contrast to that paper, we rely solely on

[9] We then conduct a thorough out-of-sample evaluation of interval forecasts using all of the tests prescribed by Christoffersen (1998), which confirms that the in-sample predictive power of the estimated quantile model translates to out-of-sample performance.

[10] Kelly and Jiang (2014) also use a large cross-section of observed returns to compute a tail risk measure

returns on a broad market index and we use high-frequency intra-day data to obtain more information about the tail. Moreover, while in the aforementioned paper the authors propose the risk-neutral expected shortfall as a new tail risk measure, we concentrate on measuring the tail risk premium as the difference between expected shortfalls under the risk-neutral and physical measures.[11]

Our paper is close in spirit to Weller (2019), who develops a real-time tail risk measure based on intra-day bid and ask quotes. While Weller (2019) focuses on the natural relation between tail risk and jumps, providing substantial evidence of jump realization predictability using intra-day data, we focus on the broader relation between aversion to tail risk and aggregate and cross-sectional risk premiums. Additionally, Weller (2019)'s measure considers a panel of 2800 firms for its baseline one-factor model, while our tail measure necessitates only one day of intra-daily observations on one stock index. In fact, the measure we propose can be easily applied to different markets and assets for which intra-day returns are available, avoiding the need to rely on option prices or bid-ask spreads.

Our paper also complements the findings obtained with the use of high-frequency factor models in Bollerslev et al. (2013), Bollerslev et al. (2016) and Aït-Sahalia et al. (2020). Bollerslev et al. (2013) use a large high-frequency data set on the cross-section of stock returns to measure the quantity of tail risk under the physical probability measure. In contrast, we identify a measure of the high-frequency tail risk premium that captures the wedge between the risk-neutral and physical worlds. Bollerslev et al. (2016) consider an extension of the CAPM model with separate betas for the jump component and the continuous component of the market return. They find that only the jump component beta entails significant premiums. We offer a new tail risk premium measure that is not dependent on option data

by assuming that asset return tails follow a power law. With particular emphasis on the financial sector, Allen et al. (2012) and Brownlees and Engle (2017) adopted VaR and expected shortfall measures to estimate systemic risks. In this literature, the estimated tail risk measures are calculated on a monthly or weekly basis rather than daily.

[11]In addition, by focusing on high-frequency individual returns instead of cross-sectional returns, supplementary economic restrictions like the non-negativity of the equity premium can be naturally imposed on the Euler equations as advocated by Campbell and Thompson (2008) and Pettenuzzo et al. (2014).

and does not rely on any parametric dynamics for the market return. Aït-Sahalia et al. (2020) estimate a multi-factor model at high-frequency. They show that a large part of the market equity premium is due to exposures to the market's jump risk component and that jump risks in Fama-French factors supersede their continuous counterparts. The predictive power of our tail risk premium for the expected returns of the characteristic-based portfolios is consistent with their findings.

2 The Nonparametric Tail Risk Premium

2.1 Background

Let $(\Omega, \mathbb{F}, \mathbb{P})$ be a probability space (with $\mathbb{P}$ the physical probability measure), where R and R_F are random variables denoting, respectively, the return of a primitive basis asset (the stock index), and a risk-free rate. An admissible risk-neutral distribution (RND) $\mathbb{Q}$ is represented by a density q, a non-negative random variable with unitary mean satisfying the Euler pricing equation for the index returns:[12]

$$\mathbb{E}^{\mathbb{Q}}[R - R_F] \equiv \mathbb{E}^{\mathbb{P}}[q(R - R_F)] = 0, \tag{1}$$
$$\text{with } \mathbb{E}^{\mathbb{P}}[q] = 1,\, q \geq 0.$$

At each day t, we observe a sample $\{R_i^t\}_{i=1,...,T}$ of high-frequency stock index returns $(T > 1)$. We use this high-frequency time-series to identify a conditional RND $\mathbb{Q}_t$ via its density q^t and the empirical conditional physical distribution $\mathbb{P}_t$, with density $p_i^t = \frac{1}{T}, i = 1, ..., T$, for all t.[13]

[12]A RND is a probability distribution $\mathbb{Q}$ equivalent to the physical distribution $\mathbb{P}$, under which the basis assets are correctly priced, i.e, it satisfies the Euler equation. It can be represented one-to-one with its corresponding density q. In the paper, we use these two definitions interchangeably.

[13]We assume stationarity and ergodicity of the composite process $(q_i^t, R_i^t)_{\{i=1,...T\}}$, such that it satisfies a time-series version of the law of large numbers (Hansen and Richard, 1987).

2.2 Definition

Our objective is to build a simple tail risk premium measure that depends solely on the returns of the single stock index observed at a high frequency. To that end, for each date t, we estimate daily expected shortfalls under the conditional physical and risk-neutral measures:

$$ES_t^{\mathbb{P}_t} := \mathbb{E}^{\mathbb{P}_t}[(s_\alpha - R)^+], \tag{2}$$

$$ES_t^{\mathbb{Q}_t} := \mathbb{E}^{\mathbb{Q}_t}[(s_\alpha - R)^+]. \tag{3}$$

In the equation above, t is the estimation date, α is a confidence level, s_α is the α-quantile of R under the physical probability, and $\mathbb{P}_t$ and $\mathbb{Q}_t$ indicate the physical and risk-neutral conditional probability densities at time t, respectively. We adopt ES to measure risk since it is a coherent measure of risk (Artzner et al., 1999), which overcomes the main deficiencies of the Value-at-Risk (VaR) measure. In particular, while VaR completely ignores the behavior of returns in the tail beyond its confidence level, ES takes an average of these tail returns, being thus highly sensitive to what happens in the tail.[14]

Note that we use a version of ES that can also be interpreted as the expectation of the payoff $(s_\alpha - R)^+$ of a put option with strike s_α. $ES_t^{\mathbb{Q}_t}$ is then the option price computed under the risk-neutral distribution incorporating investors' preferences. The put payoff is positive in the states of nature where $R < s_\alpha$, which for $\alpha < 0.5$, at least for symmetric distributions, occurs for returns R smaller than the risk-free rate R_F, i.e., negative excess returns. The smaller the α, the more negative the s_α and the farther we are in the left tail of R. To strike a good balance between capturing behavior in the left tail and guaranteeing enough return observations to extract information from, we set $\alpha = 0.2$.[15] This essentially means that $(s_\alpha - R)^+$ is the payoff of an OTM put option, whose price $ES_t^{\mathbb{Q}_t}$ reflects how protection against downside risk (the risk of large negative returns below s_α) is valued according to $\mathbb{Q}_t$.

[14] See Berkowitz and O'Brien (2002) and Jorion (2019) for examples in which VaR leads to overestimation of risks in calm periods, and underestimation during crises.

[15] In Section 4.3, we show that our results are very similar if we define α to be 0.1 instead.

We define our tail risk premium measure as the difference between the expected shortfalls under the risk-neutral and physical distributions:

$$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t = ES_t^{\mathbb{Q}_t} - ES_t^{\mathbb{P}_t}. \tag{4}$$

Given a positive equity premium, negative market return states in which the OTM put pays off are deemed more likely to happen under $\mathbb{Q}_t$ than under $\mathbb{P}_t$ due to risk aversion. By keeping the same threshold s_α for the two ES's, the tail risk premium depends only on how investors' preferences encoded in risk-neutral probabilities make $ES_t^{\mathbb{Q}_t}$ exceed its physical counterpart $ES_t^{\mathbb{P}_t}$. In particular, $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t$ can be interpreted as the expected gain of selling the put option by noting that $ES_t^{\mathbb{Q}_t} - ES_t^{\mathbb{P}_t} = \mathbb{E}^{\mathbb{P}_t}[\mathbb{E}^{\mathbb{Q}_t}[(s_\alpha - R)^+] - (s_\alpha - R)^+]$ is the expected value of the difference between the put price and its payoff. Therefore, $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t$ is expected to be high when investors highly value protection against market downside risk (i.e., when they are willing to pay a high premium relative to the expected put payoff).[16]

In our empirical application, we compute $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t$ at a daily frequency using the intra-day stock index returns. Since there is no overlapping of data when calculating our measure, it avoids spurious persistence and is able to quickly react to the arrival of new information embedded in market returns. At this point, our main challenge lies in finding a way of identifying the conditional risk-neutral probability distribution $\mathbb{Q}_t$ without using options in the estimation process. Following Almeida et al. (2017), who compute risk-neutral probability distributions from a cross-section of portfolio returns, we obtain $\mathbb{Q}_t$ at each date t by solving a specific minimum distance problem between the conditional physical probability distribution $\mathbb{P}_t$ and the family of risk neutral distributions that correctly price the S&P 500 returns $\{R_i^t\}_{i=1,\ldots T}$ within day t. The details of this procedure are developed in the next subsection.

[16] In Appendix A, we also show that our tail premium measure captures time-variation in the risk-neutral and physical tail shape parameters of the market return distribution.

2.3 Identifying Conditional Risk-Neutral Densities from S&P 500 Returns

We work with a sequence of repeated one-period models indexed by time t, characterized by the high-frequency returns $\{R_i^t\}_{i=1,...,T}$ sampled at t.[17] Since $T > 1$, the market within each of these models is inherently incomplete and, under the assumption of no-arbitrage, there exists an infinity of RNDs pricing the index returns at each date t. Almeida and Garcia (2017) suggest identifying a subset of RNDs by minimizing functions in the Cressie-Read family of discrepancies $\phi^\gamma(\pi) = \frac{\pi^{\gamma+1}-1}{\gamma(\gamma+1)}, \gamma \in \mathbb{R}$, that measure the distance between admissible RNDs $\mathbb{Q}$ with density q and the physical probability distribution $\mathbb{P}$ with density p. Each discrepancy $\phi^\gamma(\pi)$ allows for the identification of a specific RND $\hat{q}^\gamma$ with unique sensitivity to higher-order moments of the stock index returns.[18]

The minimum-discrepancy (MD) problem, which was originally proposed by Almeida and Garcia (2017) for SDFs and later adapted by Almeida and Freire (2022) for RNDs, can be stated in its sample version for a single basis asset as follows:

$$\begin{aligned} \hat{q}^\gamma = \underset{\{q_1,...,q_T\}}{\arg\min} \quad & \sum_{i=1}^{T} p_i \phi^\gamma\left(\frac{q_i}{p_i}\right) \\ \text{subject to} \quad & \sum_{i=1}^{T} q_i (R_i^t - R_F) = 0 \\ & \sum_{i=1}^{T} q_i = 1 \\ & q_i \geq 0 \,(\text{or } q_i > 0) \; \forall \, i, \end{aligned} \tag{5}$$

where the last inequality depends on the discrepancy $\phi^\gamma(.)$ chosen to measure the distance between the RND q and the physical density p: if $\gamma > 0$, then $q \geq 0$, otherwise $q > 0$. We consider homogeneous empirical probabilities $p_i = \frac{1}{T}, i = 1, 2, ..., T$, to represent the physical

[17] We often omit the time t dependence for ease of notation.

[18] The original analysis in Almeida and Garcia (2017) is performed in terms of Stochastic Discount Factors (SDFs). Almeida and Freire (2022) adapt it in details to consider RNDs instead. Most of the technical details following below were originally derived in these two papers.

distribution, for all dates t.[19] This allows us to exchange the expectation $\mathbb{E}^{\mathbb{P}}$ with its sample counterpart $\frac{1}{T}\sum_{i=1}^{T} \equiv \sum_{i=1}^{T} p_i$.

The MD problem (5) is computationally simpler and faster to solve in its dual formulation:

$$\hat{\lambda}^{\gamma} = \arg\sup_{\theta \in \mathbb{R},\, \lambda \in \Lambda} \theta + \sum_{i=1}^{T} p_i \phi^{*,\gamma}(\theta + \lambda(R_i^t - R_F)), \tag{6}$$

where $\Lambda \subseteq \mathbb{R}$ and $\phi^{*,\gamma}$ denotes the convex conjugate of ϕ^{γ}, restricted to a subset of the non-negative real line:

$$\phi^{*,\gamma}(z) = \underset{w \in [0,\infty) \cap \text{domain } \phi^{\gamma}}{sup} zw - \phi(w). \tag{7}$$

In this dual problem, θ and λ are Lagrange multipliers arising from the restrictions defining an admissible RND in (5). The multiplier θ determining that q sums to one (i.e., that q is a discrete probability distribution) can be concentrated out of the problem. On the other hand, the multiplier λ enforcing the Euler equation for the returns is the main one completely characterizing the RND. More specifically, for $\gamma < 0$, we can solve the following dual optimization problem to obtain λ:

$$\hat{\lambda}^{\gamma} = \arg\sup_{\lambda \in \Lambda} \; -\sum_{i=1}^{T} p_i \frac{1}{(\gamma+1)} \left(1 + \gamma\lambda\left(R_i^t - R_F\right)\right)^{\frac{\gamma+1}{\gamma}}, \tag{8}$$

where q can be recovered from the first-order condition of (8) with respect to λ:

$$\hat{q}_i^{\gamma} = \frac{\left(1 + \gamma\hat{\lambda}^{\gamma}\left(R_i^t - R_F\right)\right)^{\frac{1}{\gamma}}}{\sum_{i=1}^{T} p_i \left(1 + \gamma\hat{\lambda}^{\gamma}\left(R_i^t - R_F\right)\right)^{\frac{1}{\gamma}}}. \tag{9}$$

For each γ, the resulting MD RND is different. To help illustrate how $\hat{q}^{\gamma}$ depends on γ,

[19] For a set of empirically observed returns $\{R_i^t\}_{i=1,\dots T}$, where each R_i^t is independent and identically distributed according to $\mathbb{P}$, $p_i = \frac{1}{T}$ is an optimal nonparametric estimator for the physical density p (or, equivalently, the empirical measure $\mathbb{P}_T = \frac{1}{T}\sum_{i=1}^{T} \delta_{R_i^t}$ is an optimal nonparametric estimator for $\mathbb{P}$, where $\delta_{R_i^t}$ denotes a unit mass at R_i^t). See Kitamura (2006) for more details. This essentially amounts to using the histogram of returns as the empirical physical measure. Alternatively, one could use a kernel density estimator to smooth the histogram and obtain the physical probabilities. This would again amount to setting $p_i' = \frac{1}{T'}$ but for T' returns drawn from the estimated kernel density, which does not bring any further insights.

we replicate the Taylor expansion of the expected value of the discrepancy $\phi(\pi) = \frac{\pi^{\gamma+1}-1}{\gamma(\gamma+1)}$ around 1 executed in Almeida and Freire (2022):

$$\mathbb{E}(\phi^{\gamma}(\pi)) = \frac{1}{2}\mathbb{E}(\pi-1)^2 + \frac{\gamma-1}{3!}\mathbb{E}(\pi-1)^3 + \frac{(\gamma-1)(\gamma-2)}{4!}\mathbb{E}(\pi-1)^4 + ... \qquad (10)$$

As can be seen, minimizing the expected value of a given discrepancy amounts to minimizing a particular combination of higher-order moments of the RND determined by γ. In particular, for $\gamma < 1$, the weight given by the discrepancy to the skewness of the RND is negative, such that skewness is maximized, while the weight assigned to kurtosis is positive, such that kurtosis is minimized. This is consistent with a "preference" for positive skewness and "aversion" to kurtosis of stock index returns, characteristics which are in line with the findings of Kraus and Litzenberger (1976) and Backus et al. (2011). The lower the γ, the higher is the relative importance of skewness and kurtosis in the estimation of the RND.

Since Cressie Read discrepancies indexed by negative γs have higher sensitivity to higher moments of basis assets returns, they constitute an adequate choice for our goal of identifying an RND to calculate the tail risk premium. Almeida et al. (2017) work with the Hellinger discrepancy ($\gamma = -\frac{1}{2}$) to identify the RND from a cross-section of stock returns, due to the robustness of this estimator reported in former econometric studies. We choose instead to conduct our baseline analysis using the $\gamma = -3$ estimator since it is more sensitive to the skewness and the kurtosis of returns.[20] In Section 4.1, we show that our results are robust to using the $\gamma = -\frac{1}{2}$ estimator.

In the next subsection, we associate the dual problem (8) to an optimal portfolio problem where a marginal investor in the S&P 500 market selects portfolio weights on the risk-free rate and the risky asset. This provides further interpretation to our approach.

[20]Even though γ can in principle attain values in the whole real line, Almeida and Freire (2022) show that the constrained optimization in the dual problem may not have a solution for extreme negative γs (usually below -5). To guarantee an admissible RND with high sensitivity to downside risk for which a solution exists, we choose to work with $\gamma = -3$.

2.4 Portfolio Interpretation

Consider a standard optimal portfolio problem for an investor with utility function within the hyperbolic absolute risk aversion (HARA) class:

$$u^{\gamma}(W) = -\frac{1}{\gamma+1}(b - a\gamma W)^{\frac{\gamma+1}{\gamma}}, \tag{11}$$

where $a > 0$ and $b - a\gamma W > 0$, which guarantees that the function u^{γ} is well-defined, concave and strictly increasing. The investor chooses how to allocate initial wealth W_0 by investing $\tilde{\lambda}$ units of wealth on the risky asset R and the remaining $W_0 - \tilde{\lambda}$ in a risk-free asset paying R_F. The optimal allocation is such that the expected utility of end-of-period wealth $W(\tilde{\lambda}) = W_0 R_F + \tilde{\lambda}(R - R_F)$ is maximized:

$$\tilde{\lambda}^{\gamma} = \underset{\tilde{\lambda}\in\mathbb{R}}{max}\ \mathbb{E}\left[u^{\gamma}(W(\tilde{\lambda}))\right]. \tag{12}$$

Almeida and Freire (2022) show that there is a one-to-one mapping between problem (12) and the population version of (8) for a given γ. This can be easily seen via the first-order condition of (12):

$$\begin{aligned} &\mathbb{E}\left\{(R - R_F)\left[b - a\gamma(W_0 R_f + \tilde{\lambda}^{\gamma}(R - R_F))\right]^{\frac{1}{\gamma}}\right\} = 0 \\ \iff &\mathbb{E}\left\{(R - R_F)\left[(b - a\gamma W_0 R_f)(1 + \gamma\hat{\lambda}^{\gamma}(R - R_F))\right]^{\frac{1}{\gamma}}\right\} = 0 \\ \iff &(b - a\gamma W_0 R_f)^{\frac{1}{\gamma}}\mathbb{E}\left\{(R - R_F)(1 + \gamma\hat{\lambda}^{\gamma}(R - R_F))^{\frac{1}{\gamma}}\right\} = 0 \\ \iff &\mathbb{E}\left\{(R - R_F)(1 + \gamma\hat{\lambda}^{\gamma}(R - R_F))^{\frac{1}{\gamma}}\right\} = 0, \end{aligned} \tag{13}$$

where $\hat{\lambda}^{\gamma} = -\tilde{\lambda}^{\gamma} a/(b - a\gamma W_0 R_f)$. The above shows that the RND $\hat{q}^{\gamma}$ in (9) is proportional to the marginal utility of the HARA investor with concavity parameter γ, and that the optimal Lagrange multiplier $\hat{\lambda}^{\gamma}$ is proportional (with opposite sign) to the optimal portfolio weight $\tilde{\lambda}^{\gamma}$ in the risky asset.

Provided that the equity premium is positive, i.e., $\mathbb{E}(R - R_F) > 0$, the optimal portfolio solution will contemplate buying a certain amount of the stock index (that is, $\tilde{\lambda}^\gamma > 0$). In such case, the Lagrange multiplier $\hat{\lambda}^\gamma$ is negative and the RND $\hat{q}^\gamma$ will distort the original physical distribution p by putting higher (lower) probability mass to any state of negative (positive) index excess return. Intuitively, the marginal utility of the HARA investor is high (low) for negative (positive) realizations of the optimal portfolio due to risk aversion. Since negative return states are the ones that matter for the expected shortfall, this means that the risk-neutral leg of our tail risk premium measure, $ES^{\mathbb{Q}}$, is greater than the corresponding physical leg, $ES^{\mathbb{P}}$, implying that $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t \geq 0$.

In practice, on a given day t, it may happen that the sample average of the high-frequency excess market returns is negative. This would imply that the HARA investor short-sells the stock index and has low marginal utility for negative market return states, such that $ES^{\mathbb{Q}}$ would be actually below $ES^{\mathbb{P}}$. Arguably, this is not economically sound, as the marginal investor selling the market would be inconsistent with a representative agent equilibrium model. In fact, the average intra-day return on day t is a noisy estimate of the conditional equity premium, which is often considered to be non-negative (Campbell and Thompson, 2008; Martin, 2017). Therefore, in our baseline analysis, we restrict the equity premium to be non-negative, which guarantees that our tail risk premium will be a non-negative measure.[21] In Section 4.2, we consider the unrestricted case for robustness and show that this restriction has no material impact on our empirical results.

3 Empirical Analysis

Our empirical analysis of the predictability of risk premiums associated with the U.S. aggregate market, identified by the S&P 500 index, is composed of three parts. The first

[21]That is, for each day t, if the average of the intra-day excess market returns $\{R_i^t\}_{i=1,\dots,T}$ is negative, we shift the mean of the return distribution to zero by transforming returns to $\tilde{R}_i^t = R_i^t - \frac{1}{T}\sum_{i=1}^{T} R_i^t$. Note that this transformation does not affect the higher (centered) moments of the return distribution.

part focuses on the predictive power of our tail risk premium ($\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$) and tail risk realization ($ES^{\mathbb{P}}$) measures for the equity premium and the variance risk premium. In part two, we extend the study to the cross-sectional predictive ability of our measures. The third part examines the broader implications for the predictability of different quantiles of the distribution of excess market returns.

3.1 Data Description

Our dataset is compiled from a number of data sources and covers the period from January 2004 to December 2018. First, we obtain high-frequency data on the S&P 500 index from www.tickdata.com, and down-sample it to the five-minute frequency by registering the last observation in each five-minute window. This data is used not only to estimate $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and $ES^{\mathbb{P}}$, but also a number of high-frequency return variation measures, such as realized variance (RV), integrated quadratic variation (IV) and jump quadratic variation (JV).[22] Data on the VIX index is obtained from the CBOE via WRDS, whereas the data on daily close-to-close S&P 500 returns, inclusive of dividends, is obtained from CRSP. Similarly to Bollerslev et al. (2009), we define the variance risk premium as the RV minus the VIX squared. The risk-free rate and data on the cross-sectional stock portfolios is obtained from the Kenneth French Data Library.[23] High-frequency data on options on the S&P 500 index is sourced from the CBOE in the form of best available bid and ask price quotes at the end of every one-minute period, and then down-sampled to the five-minute frequency (the procedure to calculate option returns is described in Section 4.4).[24] Finally, the daily time-series of the Left Tail Variance (LTV) is obtained from www.tailindex.com. Table B.1 collects all the variables' definitions.

[22] RV is estimated as in Andersen et al. (2003). IV is estimated as in Mancini (2009). JV is estimated as $\max\{RV - IV, 0\}$.

[23] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[24] For each day, we use the options with maturity closest to one month, taking the shorter maturity as a tie breaker. We remove all quotes with zero bid prices and those where the ask price is more than five times the bid price.

3.2 Properties of the Tail Risk Premium and Predictors

Our tail risk premium measure $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is based on the risk adjustment provided by the risk-neutral distribution methodology described in Section 2.3. Figure 1 presents the time-series of the physical expected shortfall $ES^{\mathbb{P}}$, the tail premium $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$, and the Lagrange multiplier (LM) for the S&P500 returns coming from the dual problem solved to obtain the RND on each day. A negative LM means that the marginal investor is long in the index. The third panel in Figure 1 shows that the investor is always long in the index, which is a direct consequence of the economic restriction of a non-negative equity premium that we impose. This restriction also guarantees that our tail risk premium is always above or equal to zero.[25] The first two panels of Figure 1 further reveal that both $ES^{\mathbb{P}}$ and $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ achieve their peaks during the 2008-2009 subprime crisis. On the other hand, the marginal investor's position in the S&P 500 index, which is proportional to minus the LM, is smaller during times of crisis as compared to non-crisis periods. This is also intuitive and consistent with flight-to-safety: the agent invests less in the risky asset during crises.

[Figure 1 about here.]

Table 1 reports the persistence of the tail risk premium and the predictor variables and their correlations. Naturally, measures of risk such as $ES^{\mathbb{P}}$, LTV, RV, IV and JV are strongly persistent and can be predicted by many of their lags. This is in contrast to risk premium measures. The VRP is only significantly predicted by its first lag, whereas the tail risk premium $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ can only be predicted by the sum of lags 11 through 22. As for the correlation matrix in Panel B, it can be seen that our tail risk premium measure is not strongly correlated with any other variable. In particular, it has a 0.26 and 0.04 correlation with $ES^{\mathbb{P}}$ and VRP, respectively. This suggests that our measure captures distinct information relative to the physical expected shortfall and the variance risk premium. On the other

[25] We relax this restriction in Section 4.2 and compare predictability results. In the unrestricted case, for a significant number of days the average of intra-day S&P500 excess returns is negative. Nonetheless, this does not affect the predictive ability of our tail premium measure.

hand, $ES^{\mathbb{P}}$ has large correlations with the volatility and jump risk variables. In what follows, we test the ability of these measures to forecast risk premiums and the future distribution of S&P 500 returns.

[Table 1 about here.]

3.3 Predicting Risk premiums

3.3.1 Equity Premium

Table 2 reports the forecasting results for 1-day ahead excess returns on the S&P 500 index based on daily predictive regressions. Excess returns are obtained over the 3-month risk-free rate reported on Kenneth French's data website, appropriately pro-rated. We investigate the market return predictability afforded by our tail risk premium while controlling for several sets of predictors. More specifically, we include as controls the physical expected shortfall $ES^{\mathbb{P}}$, the variance risk premium VRP, the LTV of Bollerslev et al. (2015) and the RV (as is or decomposed into JV and IV).

The first column of Table 2 shows that the physical expected shortfall does not significantly predict 1 day-ahead excess market returns. When we include the tail risk premium in the regression in column (2), $ES^{\mathbb{P}}$ remains insignificant. In contrast, $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ has strong predictive power for the equity premium. This can be seen from a regression coefficient that is statistically significant at the 1% level and an adjusted R^2 that is relatively high for the 1-day horizon. The positive coefficient indicates that investors require a higher compensation to hold the market (i.e., a higher excess return) when the tail risk premium increases. These findings provide new high-frequency evidence that aversion to downside risk (but not downside risk itself, as captured by $ES^{\mathbb{P}}$) is an important determinant of the equity premium.

Column (3) of Table 2 further includes the VRP as control. The VRP is statistically insignificant and only marginally increases the adjusted R^2. Importantly, such inclusion does not affect the predictive power of $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ for excess market returns. Results are very similar in column (4) when the expected shortfall is removed from the regression. This suggests that our tail risk premium reflects distinct information from that contained in the VRP. In fact, the fifth column shows that if we orthogonalize $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ with respect to the VRP, it still retains most of its predictive power for the equity premium. While the variance risk premium is an important predictor of future market returns at relatively low-frequency horizons (Bollerslev et al., 2009), we find that the same is not true in a high-frequency environment.

Finally, in the last two columns, we include realized and option-implied measures of volatility and jump risk. The risk measures are all insignificant, but increase substantially the R^2. In particular, the LTV has no predictive power. This can be rationalized by the fact that it is computed from options with maturity between 6 and 31 trading days, which reflect market expectations over relatively long horizons. When these additional variables are included, the VRP becomes significant at the 5% level, with the expected sign: investors demand higher market returns when they are more averse to volatility risk (i.e., when $VRP = RV - VIX^2$ is more negative). As for our tail premium measure, it remains significant at the 1% level. In Table C.1 in Appendix C, we report the market return predictability results for 5- and 21-day horizons. For these horizons, no variable is able to significantly predict market returns. This suggests that the primary effect of $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ on market returns is concentrated on the day following a shock in tail risk premium.

[Table 2 about here.]

We also investigate whether the market return predictability afforded by $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is robust out-of-sample. To do that, we estimate 1-day ahead predictive regressions using only data prior to a given date and keep the parameters fixed to predict the excess market return on day $t + 1$ given $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ on day t. We consider different starting dates for the out-of-sample period (2008-01-01, 2012-01-01, 2015-04-08) and different update frequencies controlling how often we re-estimate the model, ranging from every month (i.e., we re-estimate the regression each month to include the most recent data) to never (i.e., we estimate the regression only once in the training sample). For comparison, we also report results for a univariate model with the VRP, which was the only relevant control in Table 2, and a bivariate model with both $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and VRP.

Table 3 contains the results in terms of the out-of-sample predictive R^2 for each model, starting date and update frequency. The univariate model based on our tail risk premium measure always generates a positive R^2, regardless of when the out-of-sample period starts and how often we re-estimate the model. This indicates that the predictive relation between $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and future excess market returns is stable and robust to different economic conditions. In contrast, the VRP leads to either small or negative R^2's. As a result, the performance of the univariate tail premium model is uniformly better than that of the bivariate model including both $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and VRP, reinforcing the superiority of our measure for predicting the equity premium.

[Table 3 about here.]

In sum, we document that our tail risk premium measure is a strong predictor of future excess market returns at the 1-day horizon. This predictability holds in- and out-of-sample and is robust to controlling for the expected shortfall, variance risk premium and several volatility and jump risk variables. Our findings are consistent with the idea that investors require a higher compensation to hold the market following an increase in aversion to downside risk as captured by $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$.

3.3.2 Variance Risk Premium

Table 4 reports the forecasting results for the 1-day ahead market variance risk premium based on daily predictive regressions. We consider the same controls as in Table 2. The first two columns show that, while downside risk as captured by $ES^{\mathbb{P}}$ has no predictive power for the variance risk premium, aversion to downside risk reflected in $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ appears as a statistically significant predictor associated with a high R^2. This provides new evidence in a high-frequency environment that part of the variance risk premium can be explained by aversion to left tail risk. As can be seen from columns (3) and (4), the lagged VRP also affords strong predictive power. Importantly, the significance at the 1% level of our tail risk premium measure is robust to the inclusion of the lagged VRP, even when we consider its orthogonalized version in the fifth column. In fact, both the $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and lagged VRP are relevant to forecast the variance risk premium, having a similar contribution to the R^2's. The last two columns reveal that the volatility and jump risk measures are insignificant and have no impact in the previous results. In Table C.2 in Appendix C, we further report the VRP predictability for 5- and 21-day horizons. Results are very similar to those of Table 4, where $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ robustly predicts the variance risk premium across virtually all horizons and regression specifications.

[Table 4 about here.]

3.3.3 The Cross-Section of Characteristic-Sorted Portfolio Returns

The previous subsections show that our tail risk premium is an important determinant of aggregate market risk premiums at high-frequency. In this subsection, we investigate whether aversion to downside risk, as captured by $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, also commands risk premium in a cross-section of characteristic-sorted portfolios. Portfolios sorted by specific characteristics reflect compensation for different types of risk beyond those related to the market portfolio. More specifically, we consider several sets of portfolios sorted according to the often-used characteristics of size, book-to-market, profitability, investment, momentum, reversal, and industry. Most of these characteristics form the basis of factor pricing models such as those of Fama and French (1993, 2015) and Hou et al. (2014).

In Figure 2, we report the results of 1-day ahead predictive regressions of characteristic-sorted portfolios. We forecast the daily excess returns for each decile portfolio obtained from sorting stocks on a given characteristic. We consider five sets of daily predictors: i) $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, ii) $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and $ES^{\mathbb{P}}$, iii) $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and RV, iv) $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and LTV, and v) $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ and VRP. For each decile portfolio, we plot a corresponding group of five t-statistics on $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$

calculated with Andrews (1991) standard errors and the adjusted R^2 of the regression. For comparison, Figure 3 reports the same results for univariate predictive regressions based on each alternative predictor ($ES^{\mathbb{P}}$, RV, LTV and VRP).

Across all sets of predictors and characteristics, the (absolute value of the) t-statistics on $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ reported in Panel (a) of Figure 2 are mostly above 2 and often above 3. This indicates that our tail risk premium measure has strong predictive power for cross-sectional risk premiums at high-frequency. Such predictive power cannot be explained by the different controls we consider. In fact, the reported t-statistics are very similar across the different sets of predictors. Panel (b) further shows that nearly all of the R^2's are between 0.5 and 1.3%, which is of the same order of magnitude of those obtained for market returns in Table 2. The alternative predictors also do not affect the R^2, with the exception of VRP which usually leads to a slight improvement. In Figure 3, we can see that the alternative predictors are individually statistically insignificant for predicting characteristic-sorted portfolios, yielding mostly negligible R^2's. This reinforces that our tail risk premium measure is the only predictor of characteristic-portfolios at high-frequency among the variables we consider.

[Figure 2 about here.]

[Figure 3 about here.]

3.4 Predicting the Distribution of Market Returns

The previous section demonstrates that it is the aversion to downside risk (as captured by $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$) rather than the level of downside risk itself (as captured by $ES^{\mathbb{P}}$) that contains relevant predictive information for risk premiums at high-frequency. In this section, we turn to an analysis of the predictive power of $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and $ES^{\mathbb{P}}$ for the whole distribution of future excess market returns in order to shed light on the differential role of these variables for explaining the equity premium.

3.4.1 Quantiles of the Distribution of Market Returns

To understand how different variables affect the distribution of future S&P 500 returns, we adopt the conditional quantile regression framework introduced by Koenker and Gilbert (1978). The conditional quantile model for S&P500 daily excess returns $\{r_t\}_{t=1,\ldots,T_q}$, reads as follows:

$$Q_{r_{t+h}}(\tau|p_t) = \theta_0(\tau) + \theta_1(\tau)\Delta^{\mathbb{P}}_{\mathbb{Q}}ES_t + \theta_2(\tau)ES_t + \theta_3(\tau)VRP_t + \theta_4(\tau)RV_t, \tag{14}$$

where $p_t = \{1, \Delta^{\mathbb{P}}_{\mathbb{Q}}ES_t, ES^{\mathbb{P}}_t, VRP_t, RV_t\}$ and the θ_js are functions mapping $\tau \in [0,1]$ into $\mathbb{R}$. This equation states that the conditional τ-quantile of the daily S&P 500 excess return distribution at time $t+h$ is a linear function of the tail risk premium and control variables at time t. In the quantile regression of r_{t+h} on the variables in p_t, the regression coefficients θ_τ are estimated by minimizing the quantile-weighted absolute value of errors:

$$\hat{\theta}_\tau = \arg\min_{\theta_\tau \in \mathbb{R}^5} \sum_{t=1}^{T_q-h} (\tau \cdot \mathbf{1}(r_{t+h} \geq p_t\theta_\tau)|r_{t+h} - p_t\theta_\tau| + (1-\tau)\cdot \mathbf{1}(r_{t+h} < p_t\theta_\tau)|r_{t+h} - p_t\theta_\tau|), \tag{15}$$

where $\mathbf{1}(.)$ is the indicator function. The predicted quantile conditional on p_t is:

$$\hat{Q}_{r_{t+h}}(\tau|p_t) = p_t\hat{\theta}_\tau. \tag{16}$$

We estimate predictive quantile regressions for the 1-day ahead ($h=1$) distribution of S&P 500 excess returns over the full sample from January 2, 2004 to December 31, 2018. The regressions are estimated for the 5th through the 95th percentiles in 10-percentage point increments, and for the median. We report the estimated coefficients and their standard deviations for all quantiles in Table 5.

We first focus on the differential contribution of $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ and $ES^{\mathbb{P}}$ for predicting the distribution of excess market returns. The statistical significance of the estimated coefficients indicates that an increase in risk ($ES^{\mathbb{P}}$) leads to a larger probability of observing both extreme negative and positive market returns, whereas an increase in the aversion to downside risk ($\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$) shifts the quantiles around the median and the whole right tail towards more positive values. These findings can be interpreted as follows. A positive shock in expected shortfall means a more volatile market, such that it is usually followed by either a large decrease or increase of the S&P 500 index. These extreme effects cancel out when predicting directly the market returns, such that $ES^{\mathbb{P}}$ is insignificant in Table 2. In contrast, a positive shock in the tail risk premium signals that investors are more averse to extreme negative outcomes, requiring a higher compensation to hold the market. This is reflected in the positive effect of $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ on essentially all quantiles of the market return distribution. Such unambiguous effect translates to the significant positive relation between our tail premium and future market returns observed in Table 2.

Table 5 also helps understand why VRP and RV lack predictive power for future excess market returns. A decrease in the VRP (i.e., an increase in the compensation required by investors to bear variance risk) leads to a positive shift in right tail quantiles, but also to a negative shift in the left tail quantiles. Similarly, RV has a positive (negative) effect on the right (left) tail quantiles. This contributes to make the predictive relation with respect to excess market returns weak, as effects cancel out. In contrast, our tail premium measure has a positive effect on all quantiles of the S&P 500 return distribution. This reinforces the predominant role that aversion to downside risk, as measured by $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$, plays on explaining future market returns at high-frequency. In the next subsection, we evaluate the quality of the estimated conditional market return distributions through an out-of-sample study of conditional interval forecasts.

[Table 5 about here.]

[Table 6 about here.]

3.4.2 Out-of-Sample Forecasts of the S&P 500 Return Distribution

To assess the out-of-sample predictive power of the different models for various parts of the return distribution, we rely on the framework developed by Christoffersen (1998) to evaluate conditional interval forecasts. Namely, we consider the likelihood ratio tests of interval forecast conditional coverage (CC), which are comprised of the joint tests of interval forecast error independence (ID) and unconditional coverage (UC). For model evaluation, we split our data into the estimation sample which contains 75% of the data (2,820 observations starting on January 2, 2004 and ending on April 7, 2015), and the evaluation sample (the remaining 940 observations ending on December 31, 2018).

Table 6 reports the results. For all intervals as previously defined, we report the p-values of the CC, ID and UC tests for several conditioning tests based on different sets of predictors. The predictor sets are: i) $ES^{\mathbb{P}}$, ii) $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, iii) $ES^{\mathbb{P}}$ and $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, iv) VRP, v) VRP and $ES^{\mathbb{P}}$, vi) VRP and $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, and vii) VRP, $ES^{\mathbb{P}}$ and $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$.[26] We also conduct a joint evaluation of interval forecasts that aggregates forecasts to six intervals spanning the quantiles (0.0 to 0.05], (0.05 to 0.25], (0.25 to 0.5], (0.5 to 0.75], (0.75 to 0.95], and (0.95 to 1.00].

For $ES^{\mathbb{P}}$, the null hypothesis of correct conditional coverage is rejected only for two intervals, while the independence of prediction errors is never rejected. For the joint interval test, CC and UC are rejected. Results for $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ are similar, but there is less efficiency for intervals covering the left tail. When $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ is added to $ES^{\mathbb{P}}$ in the column titled "Both", efficiency improves as there is no rejection in the CC joint interval test. In fact, across nearly all intervals, results are supportive of interval forecast efficiency. VRP is the predictor that leads to most rejections of the out-of-sample coverage tests. Even so, the best performance is obtained with the model including the three predictors, for which there are essentially no rejections of CC, ID and UC in the joint test and across intervals. Overall, the results support that the in-sample predictive power of the estimated quantile models containing our

[26] Due to the fact that VRP has a better performance than RV in the previous quantile regressions, we decided to drop RV. Results including RV are similar.

tail premium measure translates to out-of-sample performance.

4 Robustness Analysis

Our tail risk premium measure depends on a number of choices regarding the estimation of the risk-neutral distribution and the expected shortfall. In this section, we perturb the method to learn about its sensitivity to those choices and compare its performance to the baseline case analyzed in Section 3.

4.1 Risk-Neutralization with the Hellinger Discrepancy

Our baseline tail premium measure relies on the particular value of $\gamma = -3$ for the parameter indexing the discrepancy in the Cressie-Read family used for estimating the risk-neutral distribution. As explained in Section 2.3, this value is chosen to tilt the distribution of market returns towards large negative returns, as an investor with high aversion to downside risk does. In this subsection, we check whether this choice is essential to our findings by investigating if a less drastic risk-adjustment, implied by the Hellinger discrepancy ($\gamma = -\frac{1}{2}$), produces similar results. This discrepancy is still consistent with aversion to downside risk (i.e., large negative returns are still overweighted), but gives less probability mass to the left tail compared to $\gamma = -3$.

In Appendix D.1, we plot the time-series of the $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ obtained using the Hellinger discrepancy and the associated Lagrange multiplier in Figure D.1. The dynamics of the new $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is very similar to that obtained for $\gamma = -3$ in Figure 1. The main difference between the two is that the new $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ attains smaller values, which is natural as the expected shortfall under a risk-neutral distribution giving less probability mass to the left tail is smaller. The Lagrange multipliers, on the other hand, are larger (in absolute value) for $\gamma = -\frac{1}{2}$. This is because the investor associated with the portfolio optimization problem buys a larger amount of the risky asset (the market index) as her aversion to downside risk is smaller.

Table D.1 reproduces the main predictability analysis for the equity premium and variance risk premium using the new $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$. The results are very similar to the baseline ones for $\gamma = -3$, where the tail risk premium measure has strong predictive power across all regression specifications. This confirms that our findings are robust to a change of discrepancy measure within the Cressie-Read family.

4.2 Risk-Neutralization Without Equity Premium Restrictions

Our baseline estimation of the risk-neutral distribution imposes a non-negativity constraint for the conditional equity premium (i.e., for the average of the high-frequency excess market returns on day t). This is to prevent the marginal investor solving the dual portfolio problem from shorting the market index, which would imply that marginal utility is low for negative market return states, such that $ES^{\mathbb{Q}}$ would be below $ES^{\mathbb{P}}$. In this subsection, we drop this restriction and, after re-calculating $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$, conclude that this has no material impact on our predictability results.

In Appendix D.2, we plot the time-series of the $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ obtained without imposing the restriction on the equity premium and the associated Lagrange multiplier in Figure D.2. The immediate consequence of allowing for a negative equity premium is that the estimates of the Lagrange multiplier turn positive on a significant part of the sample (meaning that the investor sells the market index). This, in turn, makes the implied risk-neutral distribution to put higher probability weights in states of nature where the index has large positive returns (which represent negative returns of the investor's portfolio). As a consequence, ES becomes smaller under $\mathbb{Q}$ than under $\mathbb{P}$ during those dates, as visible in the central panel of Figure D.2. This difference notwithstanding, a visual comparison with our main Figure 1 uncovers similar dynamics of $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ with and without restrictions.

Table D.2 reproduces the main predictability analysis for the equity premium and variance risk premium using the $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ without restrictions. A comparison with the baseline results in Tables 2 and 4 suggests that removing the positive equity premium constraint renders the estimated $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ noisier, as the associated t-statistics and R^2s of the 1-day horizon predictive regressions become lower. Even so, the tail risk premium measure still retains statistical significance and is the strongest predictor of aggregate market risk premiums among the variables we consider. We therefore conclude that our findings are robust to removing the equity premium constraint.

4.3 Lower Expected Shortfall Threshold

In our baseline analysis, we set $\alpha = 0.2$ as the confidence level of the α-quantile of the return distribution (s_α) for calculating the expected shortfalls in equations (2) and (3). In this subsection, we examine the sensitivity of our results to this choice by setting $\alpha = 0.1$, meaning the expected shortfall threshold is farther in the left tail of the returns and there are less return observations to extract information from.

The results are collected in Appendix D.3. First, decreasing α mechanically translates to an increase in the $ES^{\mathbb{P}}$ measure, which is evident when comparing the left panels of Figures D.3 and 1. Somewhat less intuitively, we also observe changes to the estimates of the tail risk premium comparing the central panels of the aforementioned figures. With $\alpha = 0.1$, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ becomes significantly higher than with $\alpha = 0.2$ in periods of market distress. This is compatible with heterogeneous investors' attitudes towards downside risk observed under different extreme quantiles of the distribution of returns.[27]

Table D.3 contains the main predictability results for risk premiums using the expected shortfall and tail premium measures estimated with $\alpha = 0.1$. We do not observe any changes relative to our baseline analysis regarding the predictive power of $ES^{\mathbb{P}}$ and $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ for the equity and variance risk premiums. That is, aversion to downside risk (as captured by our tail premium measure) has strong predictive power, while downside risk itself (as captured by the expected shortfall) remains a statistically insignificant predictor.

[27]See Castro and Galvao (2019) for the development of a rational dynamic model based on quantile utility preferences.

4.4 Including Option Returns in the Estimation of the Risk-Neutral Distribution

Finally, our baseline specification to compute the risk-neutral expected shortfall relies on a nonparametric risk-neutral distribution extracted solely from high-frequency market returns. However, in principle the risk-neutral measure can be estimated using information from any security for which high-frequency return data is available. In particular, given that options are informative about higher-order risks of the underlying asset, a natural question is whether adding index option returns to the estimation of the risk-neutral distribution helps the resulting tail risk premium measure better predict the equity premium. We use high-frequency data on S&P 500 options to test this possibility.

Using our high-frequency option data, we calculate the index option returns as follows. In each five-minute window, we sort out-of-the-money (OTM) call (put) options into five portfolios based on the absolute value of their Black-Scholes delta, denoted as $\Delta 01$ (deep OTM options with absolute deltas ranging from 0.0 to 0.1) through $\Delta 05$ (close to at-the-money (ATM) options with absolute deltas ranging from 0.4 to 0.5). Next, we calculate the mid price for each option. In the following step, we match the observations in a given five-minute window to those in the subsequent window, and we discard options with no match. We further drop from the return calculation the options whose mid prices did not change between two observation windows. Finally, from the remaining data we calculate the equally-weighted return on each option portfolio.

We estimate the risk-neutral distribution by solving the minimum discrepancy problem in (5) including as basis assets the S&P 500 market index and one option portfolio at a time.[28] The excess market return predictability results for the tail premium are reported in Table 7, where each column corresponds to the estimation including option returns grouped by the indicated Δ. We include $ES^{\mathbb{P}}$ as control and consider the 1-day horizon. The tail risk

[28] We refrain from including all option portfolios at once in the estimation as the number of basis assets would be too large relative to the number of high-frequency observations, which can lead to unstable results.

premium $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ is statistically significant only for the call groups $\Delta 01$ and $\Delta 03$. In terms of R^2, this model performs worse than the equivalent one without options in Table 2 (0.63 and 0.65 for $\Delta 01$ and $\Delta 03$, respectively, against 0.73). For all other option portfolios, there is little predictive power coming from the associated tail risk premium.

The results above can be rationalized by the following. When we include option and index returns together in the estimation of the risk-neutral measure, we lose the interpretation of how the risk-neutral measure acts on index returns. When the S&P 500 index is the unique asset, the risk-neutral distribution overweights market negative returns and underweights positive ones, provided that the equity premium is positive. This is such that the risk-neutral expected shortfall appropriately reflects downside risk of the market index. In contrast, when the index and an option portfolio are basis assets, the risk-neutral distribution overweights negative returns coming from the optimal solution of the dual portfolio problem (i.e., a linear combination of index and option returns), which are not necessarily negative realizations of the market. Table 7 shows that this hurts performance, giving support to our baseline specification using only market returns for the estimation.

[Table 7 about here.]

5 Conclusion

In this paper, we propose a new method to compute tail risk premium at high-frequency using solely intra-day market returns and a risk-neutralization algorithm. Empirically, we show that our tail risk premium measure has strong predictive power for aggregate market and cross-sectional risk premiums at short horizons. Such predictability is robust to controlling for established measures of risk and risk premiums and to different specifications of our measure. Our findings provide new high-frequency evidence that aversion to downside risk is fundamental to explain asset pricing behavior. A natural extension of our method would be to use a cross-section of high-frequency returns to investigate to what extent this would improve the model forecasting ability.

References

Adrian, T. and M. K. Brunnermeier (2016). Covar. *American Economic Review 106*(7), 1705–41.

Ait-Sahalia, Y. and A. W. Lo (1998). Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. *The Journal of Finance 53*(2), 499–547.

Ait-Sahalia, Y. and A. W. Lo (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics 94*, 9–51.

Allen, L., T. G. Bali, and Y. Tang (2012). Does systemic risk in the financial sector predict future economic downturns? *Review of Financial Studies 25*(10), 3000–3036.

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.

Almeida, C. and G. Freire (2022). Pricing of index options in incomplete markets. *Journal of Financial Economics 144*, 174–205.

Almeida, C. and R. Garcia (2017). Economic implications of nonlinear pricing kernels. *Management Science 63*, 3361–3380.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica 71*(2), 579–625.

Andersen, T. G., N. Fusari, and V. Todorov (2017). Short-term market risks implied by weekly options. *The Journal of Finance 72*(3), 1335–1386.

Andersen, T. G., N. Fusari, and V. Todorov (2020). The pricing of tail risk and the equity premium: Evidence from international option markets. *Journal of Business & Economic Statistics 38*(3), 662–678.

Andrews, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica 59*(3), 817.

Ang, A., J. Chen, and Y. Xing (2006). Downside risk. *Review of Financial Studies 19*(4), 1191–1239.

Artzner, P., F. Delbaen, J. Eber, and D. Heath (1999). Coherent Measures of Risk. *Mathematical Finance 9*(3), 223–228.

Aït-Sahalia, Y., J. Jacod, and D. Xiu (2020). Inference on risk premia in continuous-time asset pricing models. Working paper no. 20-30, The University of Chicago Booth School of Business.

Backus, D., M. Chernov, and I. Martin (2011). Disasters implied by equity index options. *The Journal of Finance 66*(6), 1969–2011.

Bali, T. G., K. O. Demirtas, and H. Levy (2009). Is there a relation between downside risk and expected stock returns? *Journal of Financial and Quantitative Analysis 44*(4), 883–909.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), 1481–1509.

Bekaert, G. and E. Engstrom (2017). Asset return dynamics under habits and bad environment–good environment fundamentals. *Journal of Political Economy 125*(3), 713–760.

Berkowitz, J. and J. O'Brien (2002). How accurate are value-at-risk models at commercial banks? *The Journal of Finance 57*(3), 1093–1111.

Bollerslev, T., S. Z. Li, and V. Todorov (2016). Roughing up beta: Continuous vs. discontinuous betas, and the cross-section of expected stock returns. *Journal of Financial Economics 120*(3), 464–490.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *Review of Financial Studies 22*(11), 4463–92.

Bollerslev, T. and V. Todorov (2011). Tails, Fears, and Risk Premia. *Journal of Finance 66*(6), 2165–2211.

Bollerslev, T. and V. Todorov (2014). Time-varying jump tails. *Journal of Econometrics 183*(1), 168–80.

Bollerslev, T., V. Todorov, and S. Z. Li (2013). Jump tails, extreme dependencies, and the distribution of stock returns. *Journal of Econometrics 172*, 307–324.

Bollerslev, T., V. Todorov, and L. Xu (2015). Tail risk premia and return predictability. *Journal of Financial Economics 118*(1), 113–34.

Bonomo, M., R. Garcia, N. Meddahi, and R. Tédongap (2015). The long and the short of the risk-return trade-off. *Journal of Econometrics 187*(2), 580–592.

Britten-Jones, M., A. Neuberger, and I. Nolte (2011). Improved Inference in Regression with Overlapping Observations. *Journal of Business Finance & Accounting 38*(5-6), 657–683.

Brownlees, C. and R. Engle (2017). Srisk: A conditional capital shortfall measure of systemic risk. *Review of Financial Studies 30*(1), 48–79.

Campbell, J. and J. Cochrane (1999). Force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Campbell, J. Y. and S. Thompson (2008). Predicting excess stock returns out of sample: Does anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

Castro, L. d. and A. F. Galvao (2019). Dynamic quantile models of rational behavior. *Econometrica 87*, 1893–1939.

Christoffersen, P. F. (1998). Evaluating Interval Forecasts. *International Economic Review 39*(4), 841–862.

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics 7*(2), 174–196.

Drechsler, I. and A. Yaron (2011). What's vol got to do with it. *Review of Financial Studies 24*(1), 1–45.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*, 3–56.

Fama, E. F. and K. R. French (2015). Dissecting Anomalies with a Five-Factor Model. *The Review of Financial Studies 29*(1), 69–103.

Farago, A. and R. Tédongap (2018). Downside risks and the cross-section of asset returns. *Journal of Financial Economics 129*(1), 69–86.

Ghysels, E. and F. Wang (2014). Moment-implied densities: Properties and applications. *Journal of Business & Economic Statistics 32*(1), 88–111.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica 59*(3), 667–86.

Hansen, L. and S. Richard (1987). The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica 55*(3), 587–613.

Hou, K., C. Xue, and L. Zhang (2014). Digesting Anomalies: An Investment Approach. *The Review of Financial Studies 28*(3), 650–705.

Huggenberger, M., C. Zhang, and T. Zhou (2018). Forward-looking tail risk measures. Technical report, University of Mannheim.

Jorion, P. (2019). Risk management lessons from the credit crisis. *European Financial Management 15*(5), 923–933.

Kahneman, D. and A. Tversky (1979, March). Prospect Theory: An Analysis of Decision under Risk. *Econometrica 47*(2), 263–291.

Kelly, B. and H. Jiang (2014). Tail risk and asset prices. *Review of Financial Studies 27*(10), 2841–2871.

Kitamura, Y. (2006). Empirical Likelihood Methods in Econometrics: Theory and Practice. *Cowles Foundation for Research in Economics, Yale University* (1569).

Koenker, R. (2013). quantreg: Quantile Regression. R package version 5.83. *R Foundation for Statistical Computing: Vienna) Available at: http://CRAN. R-project. org/package= quantreg.*

Koenker, R. and J. Gilbert, Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Koenker, R. and J. A. Machado (1999). Goodness of Fit and Related Inference Processes for Quantile Regression. *Journal of the American Statistical Association 94*, 1296–1310.

Kraus, A. and R. H. Litzenberger (1976). Skewness preference and the valuation of risk assets. *The Journal of Finance 31*(4), 1085–1100.

Mancini, C. (2009). Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics 36*(2), 270–296.

Mancini, C. and F. Gobbi (2012). Identifying the Brownian covariation from the co-jumps given discrete observations. *Econometric Theory 28*(02), 249–273.

Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics 132*, 367–433.

McNeil, A., F. Rudiger, and E. Paul (2005). *Quantitative Risk Management*. Princeton Univ. Press.

Pettenuzzo, D., A. Timmermann, and R. Valkanov (2014). Forecasting stock returns under economic constraints. *Journal of Financial Economics 114*(3), 517–553.

Routledge, B. R. and S. E. Zin (2010, 08). Generalized Disappointment Aversion and Asset Prices. *Journal of Finance 65*(4), 1303–1332.

Siriwardane, E. (2013). The probability of rare disasters: Estimation and implications. Technical report, NYU Stern School of Business.

Vilkov, G. and Y. Xiao (2013). Option-implied information and predictability of extreme returns. Working paper, Goethe University.

Weller, B. (2019). Measuring tail risks at high-frequency. *Review of Financial Studies 32*(9), 3571–3616.

Zeileis, A., S. Köll, and N. Graham (2020). Various Versatile Variances: An Object-Oriented Implementation of Clustered Covariances in R. *Journal of Statistical Software 95*(1), 1–36.

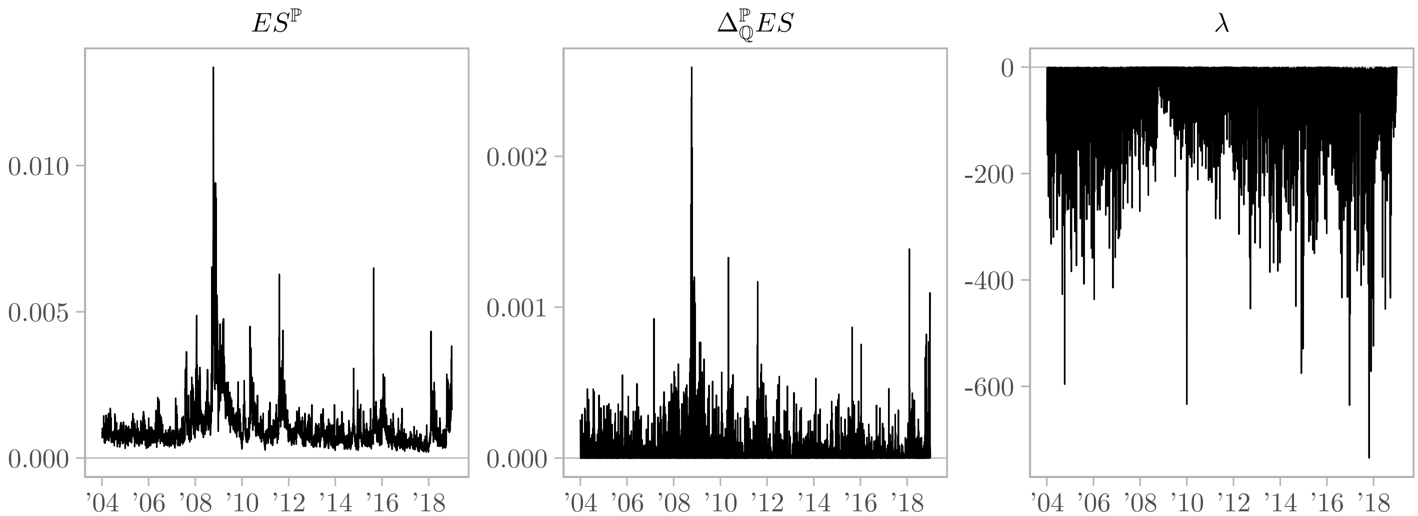

Figure 1: Time series of the tail measures implied by the S&P 500 intraday data.

Figure 2: Predictive regressions: Fama-French portfolios

Summary of predictive regressions of the 1-day ahead excess return on Fama-French portfolios of stocks sorted on Size, Book/Market, Profitability, Investment, Momentum, Reversal, and the 10 Industry portfolios. Five regression specifications are considered, using the following regressors: (i) $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$, (ii) $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and $ES^{\mathbb{P}}$, (iii) $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and RV, (iv) $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and LTV, (v) $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ and VRP. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020).

(a) t-statistics of $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$ in predictive regressions

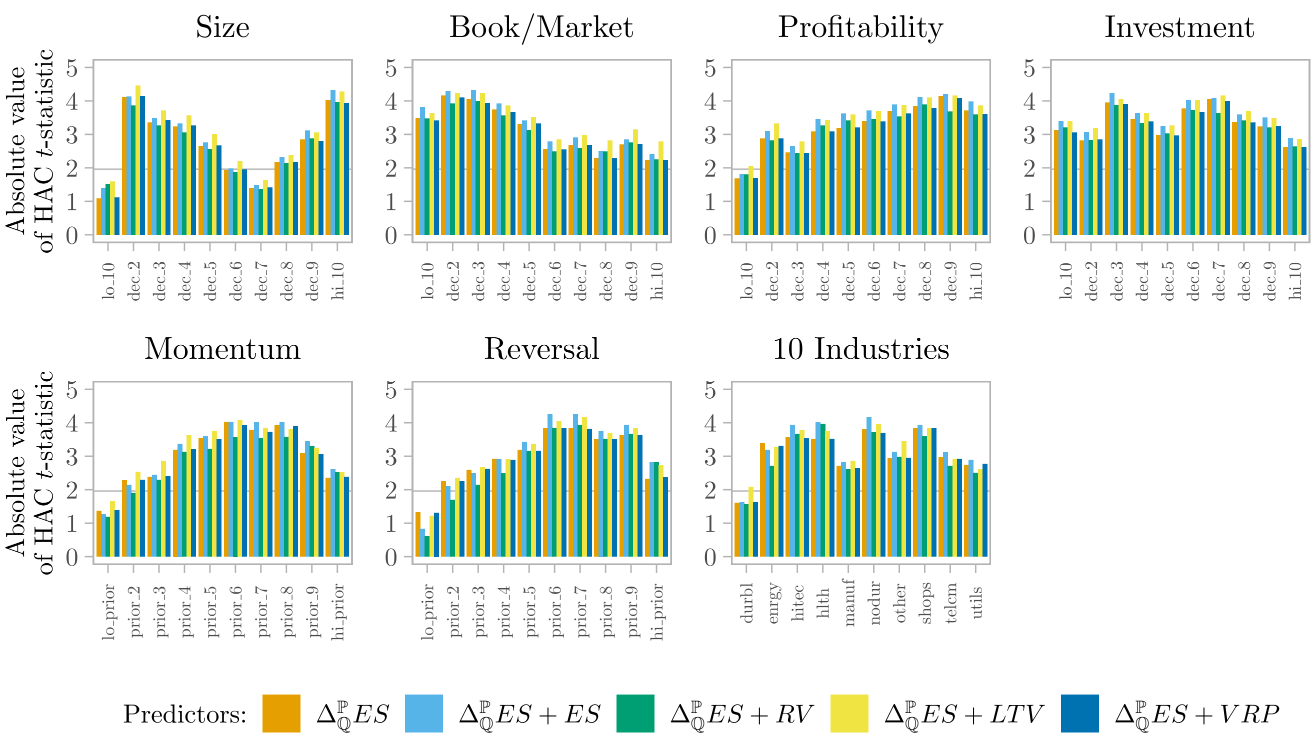

(b) R^2 of predictive regressions

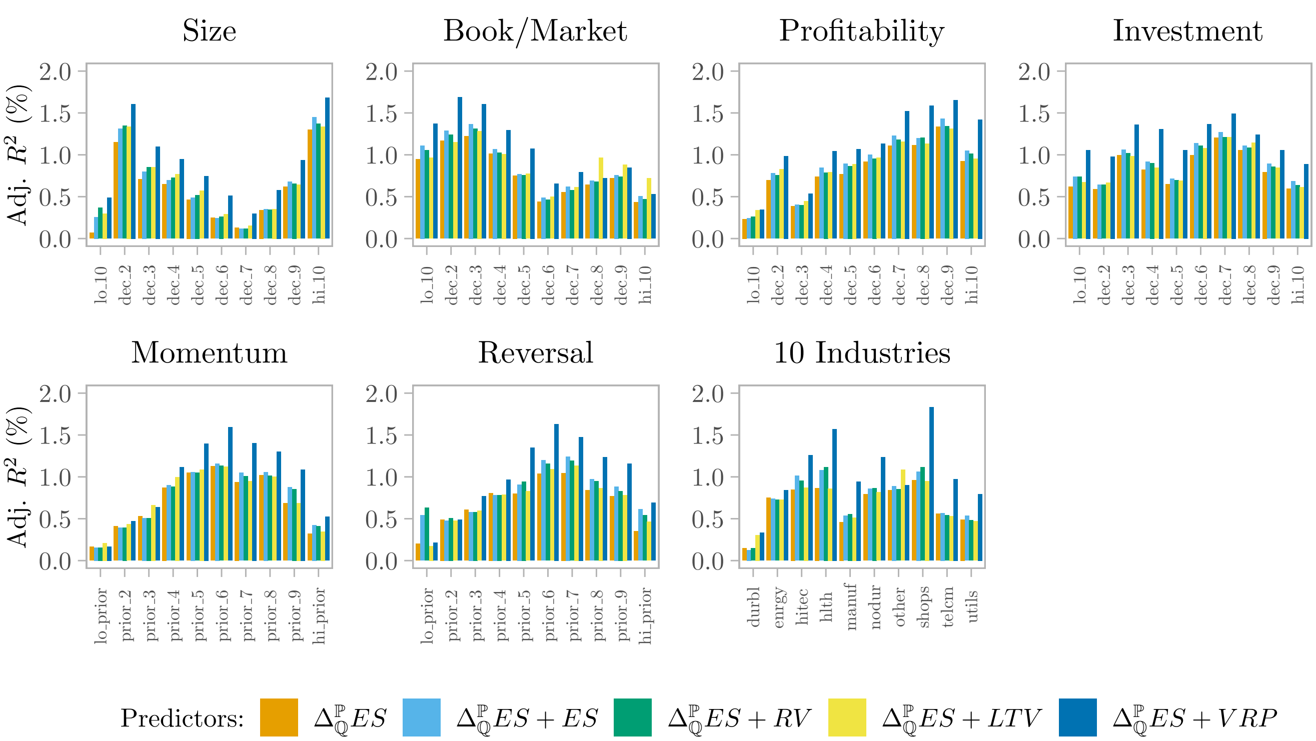

Figure 3: Predictive regressions: Fama-French portfolios, alternative predictors

Summary of predictive regressions of the 1-day ahead excess return on Fama-French portfolios of stocks sorted on Size, Book/Market, Profitability, Investment, Momentum, Reversal, and the 10 Industry portfolios. Four regression specifications are considered, using the following regressors: (i) $ES^{\mathbb{P}}$, (ii) RV, (iii) LTV, (iv) VRP. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). In order to facilitate the comparison with the main predictability results for Fama-French portfolios, the axis scales are set the same as in Figure 2.

(a) t-statistics of alternative predictors

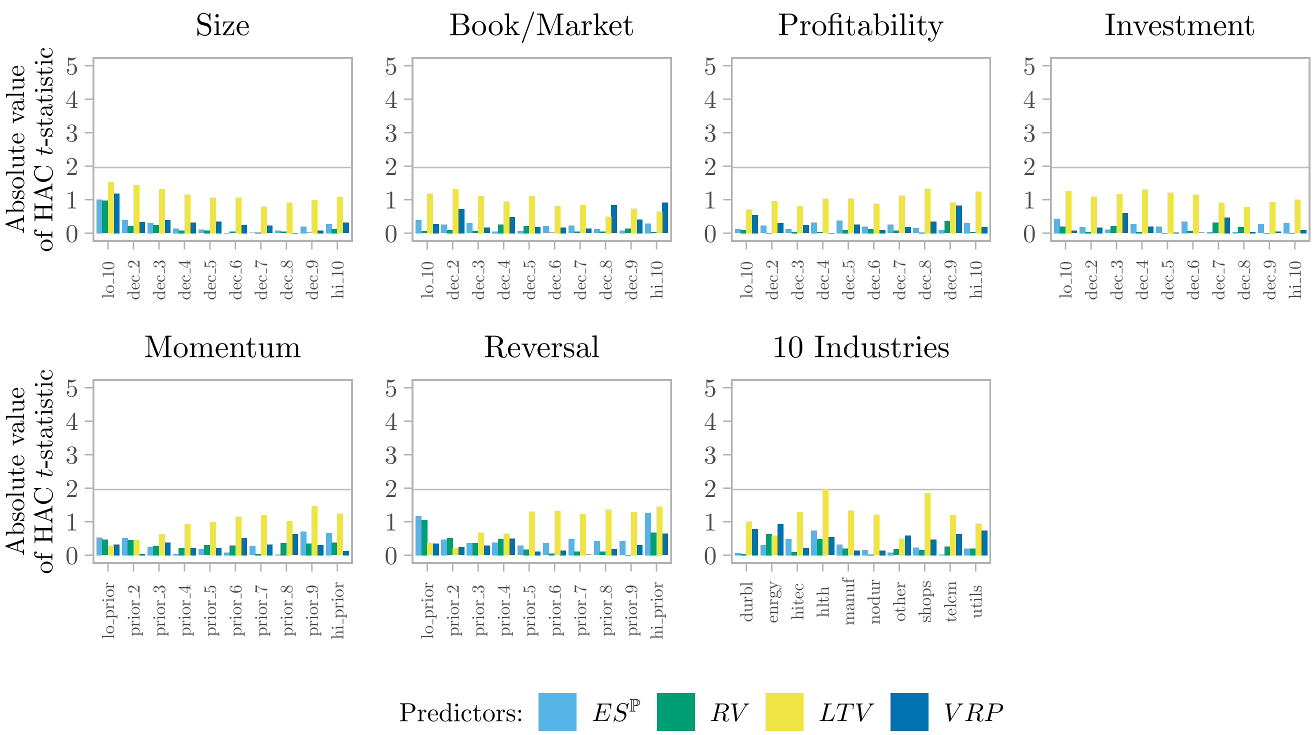

(b) R^2 of predictive regressions

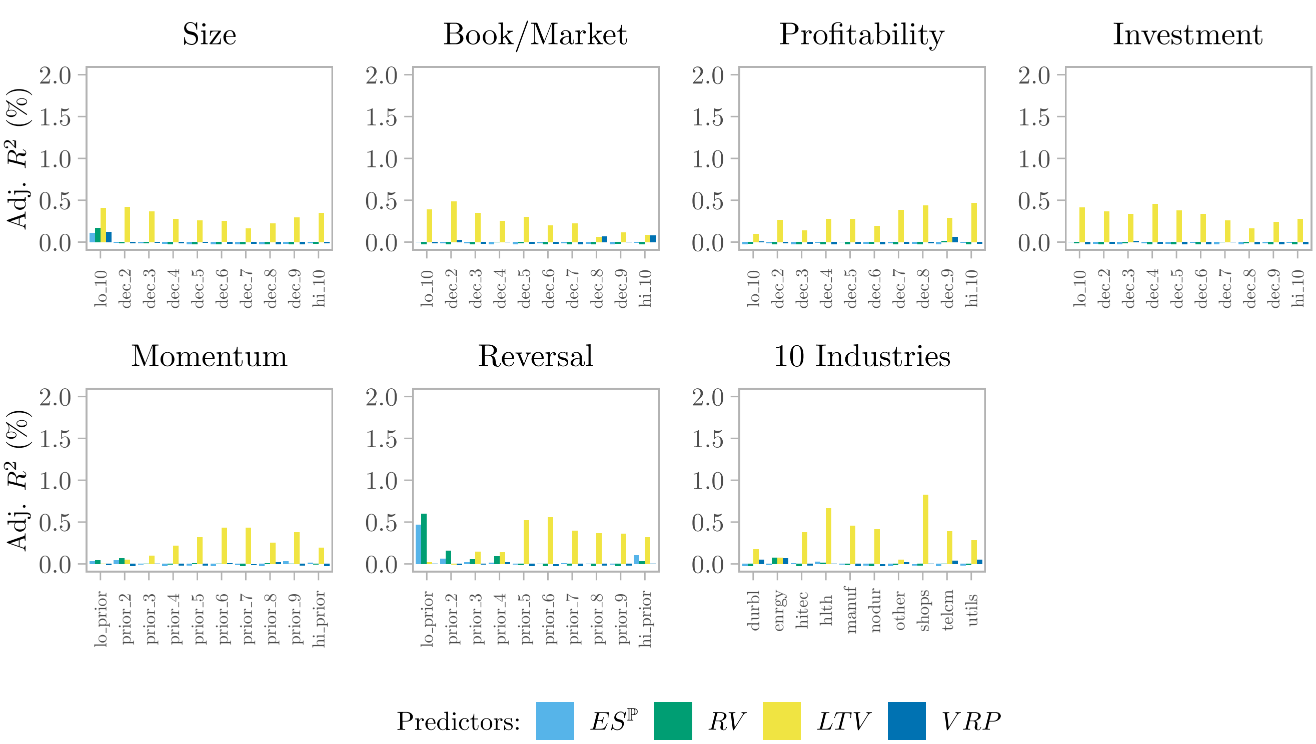

Table 1: Persistence and co-dependence of the predictors of risk premia

Panel A presents the estimates of HAR-type (see Corsi, 2009) predictive models for each risk premium predictor. The labels "Lag 1" through "Lag 5" denote the appropriately lagged regressand. The label "Lag 1W" denotes the sum of lags 6 through 10 of the regreessand, and the label "Lag 1M" the sum of lags 11 through 22. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally marked with a $\star$. We report Andrews standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3`. Panel B presents the correlation matrix of the predictors. $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES^{\mathbb{P}}_t = -E_t[R_{it}|R_{it} \leq F^{-1}_{R_t}(0.2)]$. The tail risk premium $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from `www.tailindex.com`. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

Panel A: Persistence of the Predictors of Risk Premia

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$	$ES^{\mathbb{P}}$	VRP	LTV	RV	IV	JV
Lag 1	0.00	**0.41***	**0.15***	**0.86***	**0.37***	**0.38***	**0.16***
	(0.03)	(0.04)	(0.06)	(0.05)	(0.09)	(0.10)	(0.06)
Lag 2	0.01	**0.29***	0.36	0.05	**0.42**	**0.45***	**0.22***
	(0.11)	(0.09)	(0.20)	(0.06)	(0.18)	(0.17)	(0.08)
Lag 3	−0.01	0.05	0.05	**0.28***	0.03	−0.03	**0.20***
	(0.06)	(0.07)	(0.09)	(0.10)	(0.06)	(0.07)	(0.08)
Lag 4	0.05	0.10	0.27	0.01	**0.36***	0.35	0.03
	(0.07)	(0.10)	(0.14)	(0.09)	(0.14)	(0.19)	(0.04)
Lag 5	−0.03	0.10	0.19	**0.21***	0.24	0.33	**0.14**
	(0.08)	(0.08)	(0.14)	(0.08)	(0.14)	(0.18)	(0.05)
Lag 1W	0.08	−0.03	−0.10	**−0.12***	−0.14	−0.14	**−0.12***
	(0.07)	(0.05)	(0.09)	(0.04)	(0.07)	(0.10)	(0.04)
Lag 1M	**0.02***	**0.01***	0.02	**0.01***	**0.02**	0.01	**0.06***
	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)	(0.01)	(0.02)
Constant	**0.00***	**0.00**	0.00	**0.00***	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj)	0.09	0.72	0.22	0.85	0.61	0.63	0.19
N	3748	3748	3748	3691	3748	3748	3748

Panel B: Correlation Matrix of the Predictors of Risk Premia

	$ES^{\mathbb{P}}$	VRP	LTV	RV	IV	JV
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$	0.26	0.04	0.31	0.30	0.28	0.26
$ES^{\mathbb{P}}$		0.41	0.70	0.88	0.88	0.51
VRP			0.05	0.67	0.64	0.56
LTV				0.71	0.71	0.35
RV					0.98	0.66
IV						0.51

Table 2: Predictive regressions: S&P 500 excess returns

The table reports regression coefficients and their standard deviations (in parentheses) of predictive regressions of the close-to-close S&P 500 excess returns at the 1-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES^{\mathbb{P}}_t = -E_t[R_{it}|R_{it} \leq F^{-1}_{R_t}(0.2)]$. Our Tail Risk Premium measure, $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from `www.tailindex.com`. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$ES^{\mathbb{P}}$	0.17	−0.16	−0.32			−2.04	−2.01
	(0.77)	(0.87)	(0.58)			(1.15)	(1.13)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$		**7.50***	**7.65***	**7.19***		**8.83***	**8.75***
		(2.90)	(2.63)	(2.68)		(2.83)	(2.66)
VRP			2.35	1.53	0.94	**−15.90**	**−16.05**
			(8.79)	(9.07)	(9.25)	(7.74)	(7.79)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$					**6.63**		
					(2.76)		
RV						15.25	
						(10.36)	
LTV						−0.10	−0.10
						(0.07)	(0.07)
JV							17.73
							(14.14)
IV							14.79
							(10.81)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	0.00	0.73	0.76	0.75	0.62	2.03	2.01

Table 3: Out-of-sample predictive power for S&P 500 excess returns

This table reports the out-of-sample predictive R^2 (in percent) of predictive linear models for S&P 500 excess returns estimated using two regressors: $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$ and VRP. The predictive R^2 is calculated as $1 - \frac{1}{T}\sum_t (R_t - \widehat{R}_t)^2 / \frac{1}{T}\sum_t (R_t)^2$, where R_t is the excess return on the S&P 500 index, and $\widehat{R}_t$ is a given model's prediction. Predictive regressions are estimated on expanding samples starting in January 2008 (columns (1) through (3)), January 2012 (columns (4) through (6)), and April 2015 (columns (7) through (9)). In the first two columns of each sub-panel we consider univariate predictive models, and in each respective third column, a model containing both predictors. Across the rows of the table we consider different parameter update frequencies, from every month in the top row to no updates (i.e., the parameters are estimated on data prior to the starting date and never updated) in the final row.

	OOS from 2008-01-01			OOS from 2012-01-01			OOS from 2015-04-08		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Update freq.	$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$	VRP	Both	$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$	VRP	Both	$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$	VRP	Both
1 months	0.1	−1.1	−0.3	0.8	0.0	0.7	1.5	0.0	1.4
3 months	0.0	−1.4	−0.3	0.8	0.2	0.8	1.7	0.2	1.6
6 months	0.7	0.0	0.6	0.7	0.2	0.7	1.6	0.2	1.6
12 months	0.6	−0.1	0.3	0.5	0.2	0.5	1.6	0.1	1.6
24 months	0.7	0.0	0.5	0.5	0.2	0.4	1.6	0.0	1.5
Never	0.8	0.1	0.8	0.7	0.0	0.6	1.6	−0.1	1.4

Table 4: Predictive regressions: short-term variance risk premium

The table reports regression coefficients and their standard deviations (in parentheses) for predictive regressions of the variance risk premium at the 1-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES^{\mathbb{P}}_t = -E_t[R_{it}|R_{it} \leq F^{-1}_{R_t}(0.2)]$. Our Tail Risk Premium measure, $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from www.tailindex.com. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$ES^{\mathbb{P}}$	0.02	0.00	**−0.02***			−0.03	−0.03
	(0.01)	(0.01)	(0.01)			(0.02)	(0.01)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$		**0.38***	**0.40***	**0.36***		**0.22***	**0.22***
		(0.14)	(0.14)	(0.13)		(0.06)	(0.06)
VRP			**0.32***	**0.26***	**0.23***	**0.38***	**0.39***
			(0.07)	(0.06)	(0.07)	(0.10)	(0.10)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$					**0.38***		
					(0.14)		
RV						0.01	
						(0.11)	
LTV						0.00	0.00
						(0.00)	(0.00)
JV							−0.09
							(0.26)
IV							0.01
							(0.10)
Constant	**0.00***	**0.00***	**0.00**	**0.00***	**0.00***	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	9.24	23.31	30.93	29.56	30.34	29.91	29.95

Table 5: Predictive quantile regressions: S&P 500 excess returns

The table reports quantile regression coefficients and their standard deviations (in parentheses) of quantile regressions predicting the distribution of the excess returns on the S&P 500 index at the 1-day horizon. The sample is from 2004-01-02 to 2018-12-31. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. Those significant at the 0.01 level are additionally marked with a $\star$. The standard errors are computed by pairwise bootstrap with the use of the `quantreg 5.83` (Koenker, 2013) package for `R 4.0.3`. The R^1 goodness-of-fit measure is calculated as in Koenker and Machado (1999). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES. RV is the realized variance of intra-day returns on the S&P500 Index. The daily value of the VRP is defined as $VRP \equiv RV \times 365 - (VIX/100)^2$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Quantile:	0.05	0.15	0.25	0.35	0.45	0.5	0.55	0.65	0.75	0.85	0.95
Δ_Q^PES	1.32	4.05	4.05	**5.99***	**7.39***	**7.65***	**8.05***	**7.81***	**7.41**	**10.83**	**12.27***
	(4.81)	(4.81)	(2.48)	(2.24)	(2.21)	(2.14)	(2.18)	(2.39)	(2.89)	(4.40)	(3.03)
ES	**−9.90***	**−6.40***	**−4.31***	−1.10	0.00	−0.05	0.66	1.11	**3.12***	**4.20***	**7.39***
	(2.22)	(1.65)	(1.18)	(1.17)	(0.96)	(1.02)	(0.95)	(1.04)	(1.15)	(1.18)	(1.46)
VRP	**34.43***	**19.63**	8.04	1.96	−4.94	−8.65	**−11.02**	**−14.84**	**−22.85***	**−35.73***	**−53.82***
	(11.32)	(9.21)	(7.89)	(5.81)	(5.24)	(4.66)	(5.46)	(6.65)	(8.83)	(8.10)	(7.94)
RV	−25.57	−17.39	−5.65	−10.83	−5.26	1.77	2.23	12.79	16.46	**24.11***	**31.67***
	(15.07)	(13.98)	(9.94)	(8.08)	(8.59)	(8.99)	(8.56)	(10.59)	(9.01)	(9.02)	(10.02)
Constant	**0.00**	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	**0.00**	**0.00***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^1(%)	21.55	9.36	4.04	1.45	0.56	0.63	1.09	2.85	6.16	13.45	29.82

Table 6: Out-of-sample interval forecast tests of S&P 500 excess return distribution models

The table reports the results of evaluating out-of-sample interval forecasts for the excess returns on the S&P 500 index obtained with different quantile predictive model specifications. Each row contains the *p*-values of the Christoffersen (1998) tests of interval forecast conditional coverage (CC), interval forecast error independence (ID) and unconditional coverage (UC) for the quantile interval indicated in the first column. The row denoted "Joint" contains the result of a joint evaluation of interval forecasts which follows Section 4.2 in Christoffersen (1998) and aggregates the interval forecasts to six intervals spanning the quantiles (0.0 to 0.05], (0.05 to 0.25], (0.25 to 0.5], (0.5 to 0.75], (0.75 to 0.95], and (0.95 to 1.00]. For model evaluation purposes, we split our data into the estimation sample which contains 75% of the data (2,820 observations starting on 2004-01-02 and ending on 2015-04-07), and the evaluation sample (the remaining 940 observations ending on 2018-12-31). The *p*-values which indicate the rejection of the null hypotheses at the 0.05 level are printed in bold face. Those that indicate rejections at the 0.01 level are further marked with a $\star$. All model specifications contain a constant term and the regressor set indicated on the label. $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES. The daily value of the VRP is defined as $VRP \equiv RV \times 365 - (VIX/100)^2$. "Both" indicates that the predictive model contains both the $ES^{\mathbb{P}}$ and $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ terms.

	ES			Δ_Q^PES			Both			VRP			VRP+ES			VRP+Δ_Q^PES			VRP+Both		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
Interval	CC	ID	UC	CC	ID	UC	CC	ID	UC	CC	ID	UC	CC	ID	UC	CC	ID	UC	CC	ID	UC
Joint	**0.05**	0.25	**0.01**	**0.00***	0.07	**0.00***	0.19	0.80	**0.00***	**0.00***	**0.00***	**0.00***	0.12	0.65	**0.00***	**0.00***	0.09	**0.00***	0.66	0.80	0.21
0.0 – 0.05	0.10	0.08	0.20	**0.00***	**0.02**	**0.01***	0.25	0.28	0.20	**0.00***	**0.00***	**0.02**	0.70	0.49	0.62	**0.00***	0.06	**0.00***	0.41	0.36	0.33
0.05 – 0.15	0.93	0.72	0.87	**0.00***	0.05	**0.00***	0.80	0.51	0.96	**0.00***	**0.03**	**0.00***	0.75	0.64	0.54	**0.00***	**0.02**	**0.00***	0.95	0.76	0.96
0.15 – 0.25	0.99	0.99	0.87	0.14	0.20	0.13	0.36	0.33	0.29	0.07	**0.04**	0.29	0.96	0.93	0.78	0.08	0.10	0.13	0.96	0.93	0.78
0.25 – 0.35	**0.01***	0.24	**0.00***	0.17	0.98	0.06	**0.03**	0.58	**0.01***	0.23	0.66	0.10	0.32	0.51	0.18	0.71	0.95	0.42	0.93	0.99	0.70
0.35 – 0.45	0.91	0.66	0.96	**0.00***	0.87	**0.00***	0.39	0.57	0.22	**0.00***	0.80	**0.00***	0.56	0.42	0.47	**0.00***	0.57	**0.00***	0.76	0.83	0.48
0.45 – 0.5	0.69	0.94	0.39	**0.02**	0.17	**0.01**	0.33	0.71	0.15	**0.03**	0.21	**0.02**	**0.04**	0.85	**0.01**	0.25	0.39	0.15	0.32	0.42	0.20
0.5 – 0.55	0.39	0.46	0.25	0.06	0.52	**0.02**	0.10	0.07	0.26	0.08	0.96	**0.02**	0.61	0.45	0.52	**0.00***	0.75	**0.00***	0.12	**0.05**	0.62
0.55 – 0.65	0.23	0.42	0.16	0.88	0.75	0.62	0.44	0.94	0.24	0.38	0.73	0.15	0.55	0.33	0.70	0.28	0.14	0.48	0.71	0.94	0.41
0.65 – 0.75	**0.02**	0.24	**0.01**	0.68	0.74	0.41	**0.03**	**0.02**	0.16	0.71	0.94	0.41	**0.01***	0.37	**0.00***	0.54	0.28	0.78	0.30	0.31	0.29
0.75 – 0.85	0.95	0.76	0.96	0.77	0.59	0.62	0.83	0.94	0.54	0.64	0.65	0.41	0.39	0.17	0.87	0.68	0.74	0.41	0.24	0.44	0.13
0.85 – 0.95	0.56	0.83	0.29	0.29	0.29	0.24	0.87	0.65	0.78	0.33	0.45	0.20	0.59	0.45	0.47	0.42	0.26	0.47	0.25	0.33	0.18
0.95 – 1.0	0.05	0.24	**0.03**	**0.00***	0.09	**0.00***	0.19	0.35	0.12	**0.01***	0.27	**0.00***	**0.01***	0.62	**0.00***	**0.00***	**0.02**	**0.00***	0.21	0.71	0.09

Table 7: Predictive regressions: S&P 500 excess returns with option-implied tail premium

The table reports regression coefficients and their standard deviations (in parentheses) for predictive regressions of the close-to-close S&P 500 excess returns at the 1-day horizon with $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$ calculated with the use of risk-neutral probabilities obtained from the joint risk-neutralization of high-frequency returns on the S&P 500 index and on options thereon, with options grouped into five portfolios based on their type and Δ. Δ01 corresponds to deep out of the money options while Δ05 corresponds to at the money options. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES.

	Call options					Put options				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Δ01	Δ02	Δ03	Δ04	Δ05	Δ01	Δ02	Δ03	Δ04	Δ05
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$	**6.85***	3.87	**6.09***	2.49	2.22	4.21	3.37	3.09	1.98	2.67
	(2.46)	(2.22)	(1.90)	(1.62)	(1.68)	(2.93)	(2.71)	(2.43)	(2.78)	(2.73)
$ES^{\mathbb{P}}$	−0.07	0.02	−0.04	0.10	0.04	0.00	−0.02	0.00	0.02	0.04
	(0.78)	(0.77)	(0.79)	(0.81)	(0.80)	(0.82)	(0.84)	(0.81)	(0.86)	(0.84)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	0.63	0.25	0.65	0.14	0.09	0.25	0.17	0.15	0.06	0.16

Appendix

A Extreme Value Theory, Expected Shortfall and the Tail Shape Parameter

Extreme Value Theory (EVT) shows that once we are looking far in the tail of a random variable X representing losses (negative of the returns) with a given distribution F, at values larger than an exogenous threshold u, the conditional distribution function $F(X \leq x \mid x \geq u)$ can be well approximated by a Generalized Pareto Distribution (GPD) $G(\xi, \beta)$, defined below.[29] We use this result to identify a direct link between our ES measures of risk defined under $\mathbb{P}_t$ and $\mathbb{Q}_t$ and the corresponding shape parameters $\xi_t^{\mathbb{P}}$ and $\xi_t^{\mathbb{Q}}$ that determine the thickness of the GPD's capturing the behavior of the tails of $\mathbb{P}_t$ and $\mathbb{Q}_t$.

The cumulative GPD distribution function is given by $G_{\xi,\beta}(x) = 1 - (1 + \frac{\xi x}{\beta})^{-\frac{1}{\xi}}, x \geq 0, \beta > 0, 0 < \xi < 1$, where β is a scaling parameter and ξ the shape parameter. The larger the ξ, the thicker the tail is. It is well-known that given a sample of returns $\mathcal{R} = \{R_1, ...R_T\}$ whose conditional tail distribution is approximated by a GPD and a fixed confidence level α, there is a direct relation between the expected shortfall measure with confidence level α, $ES^{\alpha}(\mathcal{R})$, the Value-at-Risk based on the same confidence level, $VaR^{\alpha}(\mathcal{R})$, and the shape parameter ξ:

$$ES^{\alpha}(\mathcal{R}) = \frac{VaR^{\alpha}(\mathcal{R})}{1-\xi} + \frac{\beta - \xi u}{1-\xi}. \tag{A.1}$$

Our tail risk measures ($ES_t^{\mathbb{P}_t}$, $ES_t^{\mathbb{Q}_t}$) and tail risk premium ($\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t$) are calculated on a daily frequency. Following Bollerslev and Todorov (2014), we allow the tail's shape parameters ($\xi_t^{\mathbb{P}}$, $\xi_t^{\mathbb{Q}}$) to be time-varying. Assuming for simplicity that the threshold u and scale β parameters are time-invariant and common to both physical and risk-neutral distributions, we can invert (A.1) to obtain the time-varying shape parameters as hyperbolic functions of our ES measures:

$$\xi_t^{\mathbb{P}} = 1 + \frac{VaR_t^{\mathbb{P},\alpha} + \beta - u}{u - ES_t^{\mathbb{P},\alpha}}, \tag{A.2}$$

$$\xi_t^{\mathbb{Q}} = 1 + \frac{VaR_t^{\mathbb{Q},\alpha} + \beta - u}{u - ES_t^{\mathbb{Q},\alpha}}. \tag{A.3}$$

Equation (A.1) also implies that our tail risk premium is a continuous function $H(.,.)$ of the time-varying tail shape parameters of both physical and risk-neutral distributions:

$$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES_t = H(\xi_t^{\mathbb{P}}, \xi_t^{\mathbb{Q}}). \tag{A.4}$$

Bollerslev and Todorov (2014) and Bollerslev et al. (2015) propose the left jump tail variance (LTV), a risk-neutral jump tail risk measure for the S&P 500 returns. They build a new dynamic model for the S&P 500 prices under the risk neutral measure $\mathbb{Q}$, whose dynamics is decomposed into a continuous stochastic volatility component and a jump component with time-varying stochastic jump-intensity. The key novelty with respect to the previous literature is the use of EVT to model the tails (both left and right) of the jump-intensity measure, suggesting that they follow Frechet distributions with time-varying shape/decay

[29] This threshold-exceedance distribution, or distribution of the tail, is usually identified based on a sample of observed random variables with the same distribution F. See Chapter 7 in McNeil et al. (2005) for more details on EVT.

parameters $F(x) = x^{-\alpha_t}, \alpha_t > 0$. They adopt the inverse of this shape parameter, $\xi_t = \frac{1}{\alpha_t}$, as a measure of tail risk. Since the Frechet distribution is directly comparable to the GPD when the shape parameter ξ is positive (i.e., $\xi_t > 0$), equations (A.3) and (A.4) directly connect our risk-neutral expected shortfall ($ES^{\mathbb{Q}}$) and tail risk premium ($\Delta^{\mathbb{P}}_{\mathbb{Q}} ES$) measures to LTV.

Moreover, since $ES^{\mathbb{Q}}_t$ is an invertible function of $\xi^{\mathbb{Q}}_t$ and $\Delta^{\mathbb{P}}_{\mathbb{Q}} ES_t$ is linear in $ES^{\mathbb{Q}}_t$, LTV and our tail measures, at least from a theoretical point of view, induce the same probability filtration (i.e., the same information sets) along the time dimension. However, they naturally differ in practice in terms of their estimation procedure and data used. LTV is estimated with the use of options. It represents the expected (risk-neutral) return volatility that stems from large negative price jumps. To identify and separate the diffusive part and the jump component of the return process, Bollerslev et al. (2015) use index options with maturity between 6 and 31 trading days, reflecting market expectations over these horizons. In contrast, our tail risk premium measure relies on high-frequency market return data, containing information completely conditional on day t.

B Variable definitions

Table B.1: Definitions of risk premiums predictors

Variable	Description
$ES^{\mathbb{P}}$	Estimate of the expected shortfall $ES^{\mathbb{P}}_t = -E_t[R_{it} \mid R_{it} \leq F^{-1}_{R_t}(0.2)]$ obtained from the empirical distribution of 5-minute returns on the S&P 500 index during market opening hours on day t. See equation (2).
$\Delta^{\mathbb{P}}_{\mathbb{Q}} ES$	Difference between $ES^{\mathbb{P}}$ and its $\mathbb{Q}$-measure counterpart; the change of measure is described in equation (1) and $\Delta^{\mathbb{P}}_{\mathbb{Q}} ES$ is defined in equation (4).
$\Delta^{\mathbb{P}}_{\mathbb{Q}} ES^{\perp}$	Residual from the full-sample regression of $\Delta^{\mathbb{P}}_{\mathbb{Q}} ES$ on VRP.
VRP	$RV - (VIX)^2/100/365$.
RV	Realized variance of 5-minute intraday returns on the S&P 500 index during market opening hours calculated as $\sum_{j=1}^{N_t} (\log r_{jt})^2$.
LTV	Left Tail Variance the S&P 500 index returns as defined in Bollerslev et al. (2015) obtained from www.tailindex.com.
IV	Integrated quadratic variation of 5-minute intraday returns on the S&P 500 index during market opening hours, the continuous component of RV, estimated as in Mancini and Gobbi (2012).
JV	The jump component of realized variance, calculated as $\max\{RV - IV, 0\}$

C Supplementary results

Table C.1: Predictive regressions: medium- and long-term S&P 500 excess returns

The table reports regression coefficients and their standard deviations (in parentheses) of predictive regressions of the close-to-close S&P 500 excess returns at the 5- and 21-day horizons. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We calculate standard errors for overlapping observations using the procedure of Britten-Jones et al. (2011). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from `www.tailindex.com`. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days
$ES^{\mathbb{P}}$	0.51	−1.66	0.03	−2.14	0.19	−0.37					−2.61	−5.10	−2.63	−4.98
	(2.88)	(10.08)	(2.77)	(9.47)	(2.64)	(9.86)					(4.48)	(13.96)	(4.41)	(13.88)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$			11.17	11.33	11.02	9.69	11.29	9.16			8.82	3.81	8.81	6.73
			(8.33)	(19.69)	(9.13)	(21.73)	(10.60)	(30.49)			(7.65)	(18.27)	(7.25)	(18.28)
VRP					−2.43	−26.50	−1.95	−27.47	−2.53	−27.73	−26.59	−82.43	−27.06	−80.94
					(16.95)	(42.86)	(17.68)	(42.73)	(17.12)	(40.36)	(25.60)	(82.05)	(25.61)	(83.73)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$									8.69	4.72				
									(10.77)	(31.82)				
RV											28.22	50.48		
											(39.78)	(122.31)		
LTV											−0.24	−0.48	−0.24	−0.49
											(0.27)	(0.70)	(0.27)	(0.70)
JV													29.57	1.10
													(45.40)	(150.93)
IV													28.25	54.15
													(40.45)	(121.18)
Constant	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.01
	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)
R^2 (adj %)	−0.01	−0.01	0.32	0.09	0.36	0.47	0.27	0.11	0.15	0.09	0.51	0.60	0.50	0.58

Table C.2: Predictive regressions: medium- and long-term variance risk premium

The table reports regression coefficients and their standard deviations (in parentheses) for predictive regressions of the variance risk premium at the 5- and 21-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We calculate standard errors for overlapping observations using the procedure of Britten-Jones et al. (2011). We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from `www.tailindex.com`. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Horizon:	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days	5 days	21 days
$ES^{\mathbb{P}}$	0.05	−0.16	0.01	−0.22	**−0.10**	**−0.57***					−0.07	−0.20	−0.07	−0.20
	(0.07)	(0.24)	(0.05)	(0.17)	(0.04)	(0.17)					(0.06)	(0.23)	(0.06)	(0.22)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$			**0.84***	**1.43***	**0.93***	**1.75***	**0.80***	0.95			**0.53***	**1.36***	**0.58***	**1.31***
			(0.27)	(0.51)	(0.23)	(0.53)	(0.26)	(0.76)			(0.12)	(0.32)	(0.10)	(0.30)
VRP					**1.59***	**5.13***	**1.34***	**3.68***	**1.26***	**3.55***	**1.69***	**6.40***	**1.73***	**6.39***
					(0.18)	(0.58)	(0.21)	(0.93)	(0.19)	(0.86)	(0.43)	(1.21)	(0.43)	(1.09)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$									**0.85***	1.20				
									(0.27)	(0.80)				
RV											−0.22	−1.95		
											(0.59)	(2.06)		
LTV											0.00	−0.01	0.00	−0.01
											(0.00)	(0.01)	(0.00)	(0.01)
JV													−1.14	−1.11
													(0.59)	(2.40)
IV													−0.14	−2.03
													(0.66)	(1.85)
Constant	**0.00***	**0.00***	**0.00***	**0.00***	**0.00**	0.00	**0.00***	**0.00***	**0.00***	**0.00***	**0.00***	**0.00***	**0.00***	**0.00***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	8.68	8.69	23.26	22.53	32.46	24.92	27.55	19.97	27.54	19.95	32.00	29.98	32.21	31.10

D Robustness results

D.1 Alternative CR discrepancy

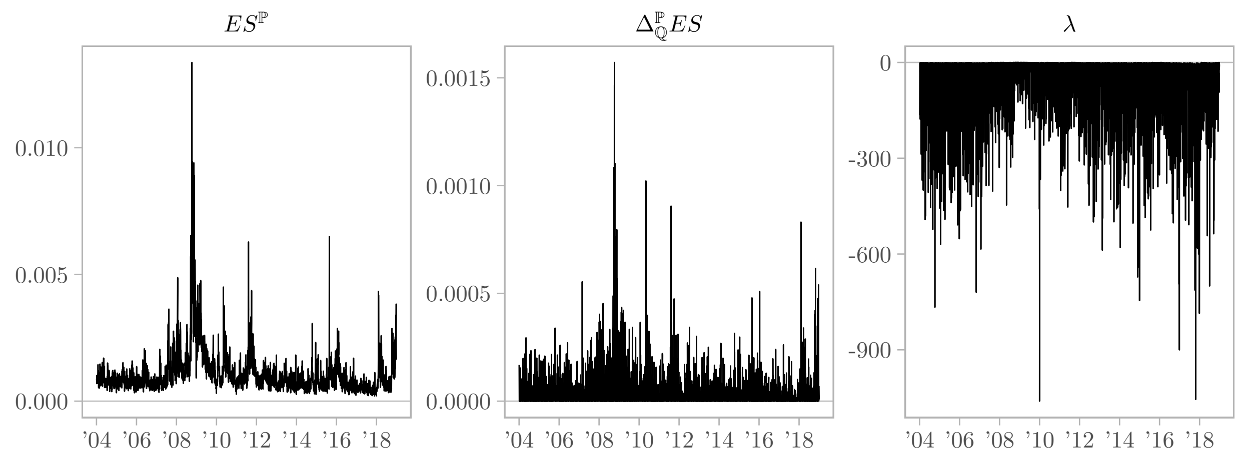

Figure D.1: Time series of the tail measures implied by the S&P 500 intraday data calculated with the Hellinger divergence.

Table D.1: Predictive regressions: S&P 500 excess returns and variance risk premium with Hellinger discrepancy

The table reports regression coefficients and their standard deviations (in parentheses) of predictive regressions of the close-to-close S&P 500 excess returns at the 1-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a $\star$. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from www.tailindex.com. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

Panel A: S&P 500 excess returns

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.17	−0.19	−0.35			−2.15	−2.12
	(0.77)	(0.87)	(0.58)			(1.14)	(1.12)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$		**11.01**	**11.26***	**10.42***		**12.57***	**12.42***
		(4.44)	(4.00)	(4.04)		(4.18)	(4.00)
VRP			2.38	1.48	0.91	**−16.18**	**−16.35**
			(8.97)	(9.06)	(9.17)	(7.79)	(7.81)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$					**9.54**		
					(4.22)		
RV						15.68	
						(10.36)	
LTV						−0.10	−0.10
						(0.07)	(0.07)
JV							18.76
							(14.39)
IV							15.16
							(10.75)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	0.00	0.67	0.70	0.68	0.56	1.95	1.93

Panel B: Variance risk premium

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.02	0.00	**−0.02***			−0.03	−0.03
	(0.01)	(0.01)	(0.01)			(0.02)	(0.02)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES$		**0.54***	**0.57***	**0.52***		**0.31***	**0.32***
		(0.20)	(0.20)	(0.20)		(0.09)	(0.09)
VRP			**0.32***	**0.26***	**0.22***	**0.37***	**0.38***
			(0.07)	(0.06)	(0.06)	(0.10)	(0.10)
$\Delta_{\mathbb{Q}}^{\mathbb{P}}ES^{\perp}$					**0.54***		
					(0.20)		
RV						0.02	
						(0.11)	
LTV						0.00	0.00
						(0.00)	(0.00)
JV							−0.06
							(0.27)
IV							0.02
							(0.10)
Constant	**0.00***	**0.00***	**0.00**	**0.00***	**0.00***	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	9.24	21.73	29.38	27.88	28.63	29.53	29.55

D.2 No equity risk premium constraints

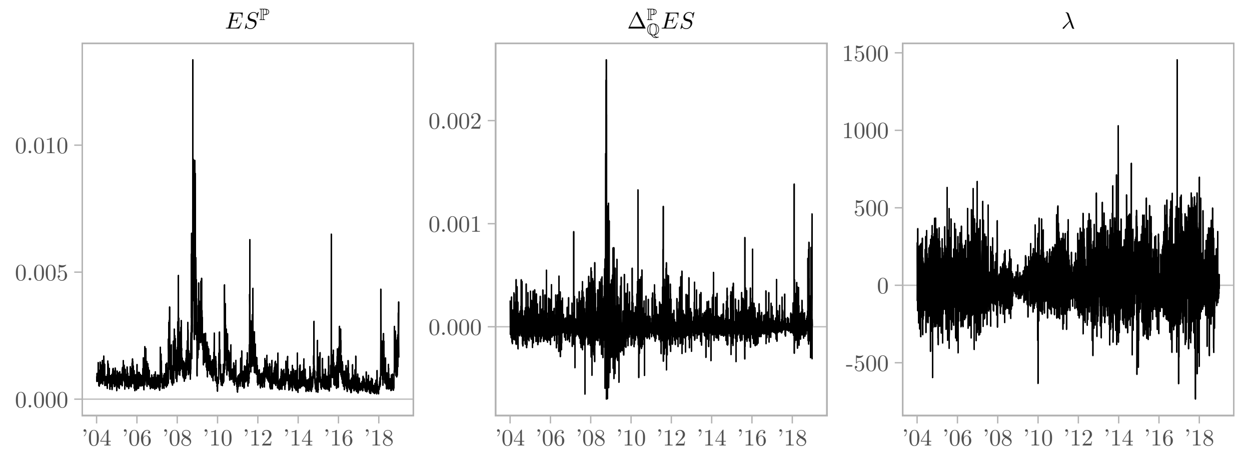

Figure D.2: Time series of the tail measures implied by the S&P 500 intraday data calculated without imposing the positive equity risk premium constraint.

Table D.2: Predictive regressions: S&P 500 excess returns and variance risk premium without imposing ex-ante equity premium constraints

The table reports regression coefficients and their standard deviations (in parentheses) of predictive regressions of the close-to-close S&P 500 excess returns at the 1-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a ⋆. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES_t^{\mathbb{P}} = -E_t[R_{it}|R_{it} \leq F_{R_t}^{-1}(0.2)]$. Our Tail Risk Premium measure, $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta_{\mathbb{Q}}^{\mathbb{P}} ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from www.tailindex.com. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

Panel A: S&P 500 excess returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.17	0.13	−0.02			−1.68	−1.65
	(0.77)	(0.84)	(0.57)			(1.21)	(1.19)
$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$		**5.44**	**5.55**	**5.55**		**6.28⋆**	**6.18⋆**
		(2.39)	(2.19)	(2.22)		(2.22)	(2.08)
VRP			2.28	2.22	1.38	**−16.05**	**−16.25**
			(8.89)	(9.01)	(9.31)	(7.72)	(7.79)
$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES^{\perp}$					**4.76**		
					(2.32)		
RV						15.23	
						(10.61)	
LTV						−0.11	−0.11
						(0.07)	(0.08)
JV							20.21
							(13.77)
IV							14.47
							(11.24)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	0.00	0.58	0.61	0.63	0.45	1.86	1.84

Panel B: Variance risk premium							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.02	0.01	−0.01			−0.02	−0.02
	(0.01)	(0.01)	(0.01)			(0.01)	(0.01)
$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES$		**0.27**	**0.28⋆**	**0.28⋆**		**0.14⋆**	**0.15⋆**
		(0.11)	(0.10)	(0.10)		(0.05)	(0.05)
VRP			**0.31⋆**	**0.30⋆**	**0.25⋆**	**0.37⋆**	**0.38⋆**
			(0.07)	(0.05)	(0.06)	(0.10)	(0.10)
$\Delta_{\mathbb{Q}}^{\mathbb{P}} ES^{\perp}$					**0.30⋆**		
					(0.10)		
RV						0.02	
						(0.11)	
LTV						0.00	0.00
						(0.00)	(0.00)
JV							−0.01
							(0.28)
IV							0.01
							(0.10)
Constant	**0.00⋆**	**0.00⋆**	**0.00**	**0.00⋆**	**0.00⋆**	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	9.24	19.99	27.40	27.29	28.20	28.22	28.19

D.3 Alternative threshold for ES calculation

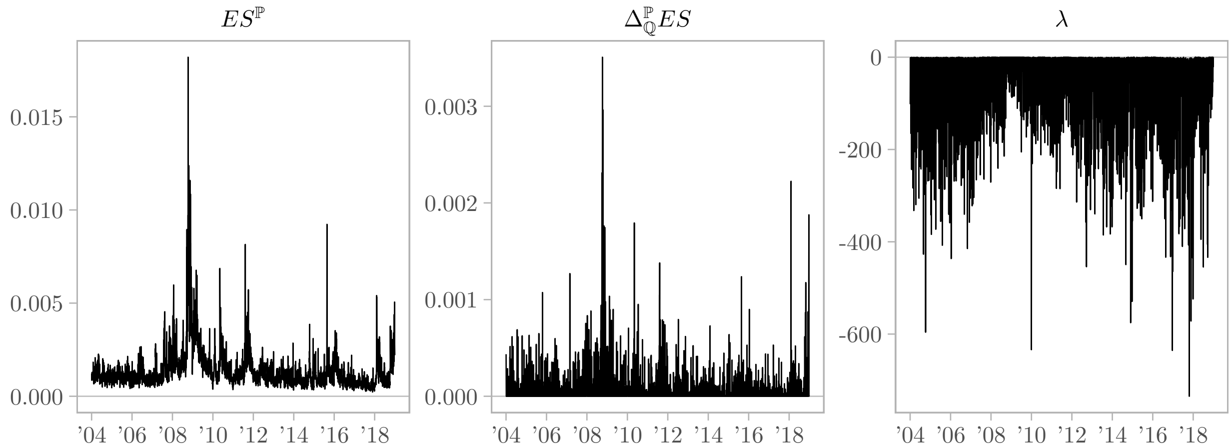

Figure D.3: Time series of the tail measures implied by the S&P 500 intraday data with $\alpha = 0.10$.

Table D.3: Predictive regressions: S&P 500 excess returns and variance risk premium with tail measures calculated at the 10th percentile of intraday returns

The table reports regression coefficients and their standard deviations (in parentheses) of predictive regressions of the close-to-close S&P 500 excess returns at the 1-day horizon. All reported coefficients and standard errors are rounded to two decimal places. Those significant at the 0.05 confidence level are printed in bold and those significant at the 0.01 level are additionally highlighted with a ⋆. We report Andrews (1991) standard errors calculated with the use of the `sandwich 3.0.0` package for `R 4.0.3` (Zeileis et al., 2020). $ES^{\mathbb{P}}$ is the non-parametric estimate of realized (physical) Expected Shortfall of intra-day S&P 500 returns before the close, on the day when each return calculation commences, $ES^{\mathbb{P}}_t = -E_t[R_{it}|R_{it} \le F^{-1}_{R_t}(0.1)]$. Our Tail Risk Premium measure, $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$, is the difference between the risk-neutral and the physical ES. VRP is the variance risk premium calculated as the difference between the day's realized variance and the (appropriately scaled) VIX^2 index. $\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$ is the component of the tail risk measure that is orthogonal to the variance risk premium. LTV is the Left Tail Variance of Bollerslev et al. (2015) obtained from `www.tailindex.com`. RV is the realized variance of intra-day returns on the S&P500 Index. IV is an estimate of integrated quadratic variation, the continuous component of realized variance, estimated as in Mancini and Gobbi (2012). JV is the jump component of realized variance, calculated as $\max\{RV - IV, 0\}$. RV, IV, and JV are calculated from return data sampled at the 5-minute frequency.

Panel A: S&P 500 excess returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.11	−0.13	−0.26			**−1.76**	**−1.75**
	(0.61)	(0.67)	(0.44)			(0.87)	(0.86)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$		**5.80⋆**	**5.91⋆**	**5.57⋆**		**6.97⋆**	**6.85⋆**
		(2.17)	(2.00)	(2.04)		(2.25)	(2.11)
VRP			2.43	1.54	1.01	**−16.39**	**−16.63**
			(8.84)	(9.14)	(9.27)	(7.48)	(7.71)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$					**5.16**		
					(2.08)		
RV						16.60	
						(9.86)	
LTV						−0.11	−0.11
						(0.07)	(0.07)
JV							20.12
							(14.14)
IV							16.14
							(10.25)
Constant	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	0.00	0.74	0.77	0.75	0.63	2.16	2.14

Panel B: Variance risk premium							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1 day	1 day	1 day	1 day	1 day	1 day	1 day
$ES^{\mathbb{P}}$	0.01	0.00	**−0.01⋆**			**−0.02**	**−0.02**
	(0.01)	(0.01)	(0.01)			(0.01)	(0.01)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES$		**0.27⋆**	**0.29⋆**	**0.27⋆**		**0.16⋆**	**0.16⋆**
		(0.10)	(0.10)	(0.10)		(0.05)	(0.05)
VRP			**0.31⋆**	**0.26⋆**	**0.23⋆**	**0.37⋆**	**0.38⋆**
			(0.07)	(0.06)	(0.06)	(0.10)	(0.10)
$\Delta^{\mathbb{P}}_{\mathbb{Q}}ES^{\perp}$					**0.28⋆**		
					(0.10)		
RV						0.03	
						(0.10)	
LTV						0.00	0.00
						(0.00)	(0.00)
JV							−0.06
							(0.25)
IV							0.03
							(0.09)
Constant	**0.00⋆**	**0.00⋆**	**0.00**	**0.00⋆**	**0.00⋆**	0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
R^2 (adj %)	9.41	22.02	29.34	28.28	28.91	29.47	29.49