Is Expected Market Volatility a Priced Risk Factor?*

Yuan Liao† Andreas Neuhierl‡ Viktor Todorov§

Abstract

We study which sources of risk in expected market volatility are priced by investors. We construct a portfolio of index options spanning short-term forward variance, whose return identifies shocks to expected volatility while remaining immune to price risk. Its average return is negative and significant. We decompose this return into three orthogonal components: dynamic market and market-neutral cross-sectional stock positions hedging variance and a stock-unrelated residual. Most of the premium comes from the dynamic market component and the residual. Thus, once dynamic market exposure is accounted for, expected market volatility is not a separate cross-sectional risk factor.

Keywords: ICAPM, market volatility, cross-sectional asset pricing, short-dated options, VIX index.

JEL classification: C51, C52, G12.

1 Introduction

The intertemporal CAPM (ICAPM) of Merton (1973) postulates that investors care about changes in the investment opportunity set, and shocks to state variables that capture these changes can therefore command risk premia in equilibrium. This insight has led to a proliferation of proposed pricing factors, see e.g., Section 9.3 in Cochrane (2005). Time-varying

*We thank Mike Chernov, Andrew Detzel, Stefano Giglio and Liuren Wu for very helpful comments. All errors are our own.

†Tippie College of Business, University of Iowa, yuan-liao@uiowa.edu

‡Mitch Daniels School of Business, Purdue University, aneuhier@purdue.edu

§Kellogg School of Management, Northwestern University; e-mail: v-todorov@northwestern.edu

expected market volatility is a leading candidate among them. In a recent contribution, Campbell et al. (2018) show that well-known cross-sectional patterns such as the value premium and the low reward to market beta can be partly understood through their exposure to market volatility risk.[1]

The central empirical challenge is that market volatility is not directly observable, and existing approaches to estimate its price of risk has well-known limitations. Parametric option-pricing studies, such as Bates (2000), Pan (2002) and Eraker et al. (2003), back out the price of volatility risk from index options under a specific model, with results that hinge critically on correct specification. Studies of delta-hedged option gains, as in Coval and Shumway (2001) and Bakshi and Kapadia (2003), rely on Black-Scholes deltas and yield positions whose returns are contaminated by price jumps. Model-free comparisons of risk-neutral and realized variance, e.g., Bollerslev et al. (2009), Carr and Wu (2009) and Drechsler and Yaron (2011), recover the realized variance risk premium, but as Bollerslev and Todorov (2011) and Drechsler and Yaron (2011) show — and as we formalize below — a large fraction of this premium is compensation for price jumps rather than for shocks to expected volatility.[2] Indeed, we show that as the realized-variance window shrinks, all of the realized variance risk is due to price jumps. Finally, the cross-sectional approach initiated by Ang et al. (2006) and pursued more recently by Detzel et al. (2023) sidesteps these issues by treating daily VIX changes as a volatility shock and estimating volatility betas from monthly regressions of daily stock returns on the market and daily VIX changes. They then form volatility-mimicking portfolios by projecting in sample daily VIX changes onto these volatility-sorted portfolios. This approach makes two (implicit) assumptions we argue are restrictive: that stocks' exposures to volatility are constant over a month, and that most priced volatility risk is spanned by the cross-section of stock returns.

Set against this background, our goal is a model-free analysis of the sources of priced risk in expected market volatility that addresses these limitations jointly. The methodological

[1]Systematic variance risk (including market variance risk) is also found to be priced in options written on the market index and individual stocks, see the recent work of Dickerson et al. (2025).

[2]Consistent with the fact that a large fraction of the realized variance risk premium is due to price jump risk, Dew-Becker et al. (2017) document initial steep increase followed by flattening of the term structure of variance swaps (whose payoff is realized variance) as their maturity increases. Dew-Becker et al. (2021) show that across a broad cross-section of options markets the premium for realized-volatility (gamma) exposure is negative while the premium for uncertainty (vega) exposure is positive, so that jumps rather than forward-looking uncertainty shocks command the risk compensation.

starting point of our analysis is the construction of a tradable portfolio of short-dated index options whose return isolates shocks to instantaneous expected market variance from price (jump) risk. The portfolio takes a long position in options expiring at horizon $2T$ and a short position in options expiring at T, so that the holding-period payoff between time s and $t < T$ spans short-term expected forward variance over $[T, 2T]$. We show that the return of this portfolio captures only shocks to expected market variance and is immune to price moves — a property that does not hold for the option portfolio behind the VIX index. We exploit two recent developments in financial markets to make this construction feasible: the rise of short-tenor index options (see Andersen et al. (2017)) and the overall increased market liquidity that allows high-frequency analysis. Because our proxy for variance shocks is the return of a traded portfolio, the price of expected variance risk is identified directly from its excess returns, with no parametric assumptions and without the need to estimate dynamic volatility risk exposures.

Our empirical application uses S&P 500 index options data from 2016 to 2025. We find a sizable and statistically significant premium for bearing expected market variance risk, with the average return on our portfolio being negative — consistent with investors disliking high-volatility states. To understand the sources of this premium, we decompose the portfolio's payoff into three orthogonal (in a martingale sense) components: a dynamic position in the market index that provides the best hedge for variance using market returns alone; a market-neutral dynamic position in the cross-section of stocks that provides the best hedge using stock-level systematic risks orthogonal to the market; and a residual component capturing variance risks that cannot be spanned by stock returns at all. We identify the systematic risks in the cross-section nonparametrically via local Principal Component Analysis, following Giglio and Xiu (2021).

The decomposition delivers a striking result. For intraday returns, essentially the entire expected variance risk premium is attributable to the dynamic market component — the market-neutral stock-spanned component earns a premium that is economically small and statistically indistinguishable from zero. This is in direct contrast to the Ang et al. (2006) finding that volatility shocks orthogonal to the market are priced in the cross-section of stocks. Once the dynamic dependence of expected variance on the market return is properly accounted for, expected market variance does not appear as a separate risk factor for stocks. Overnight, by contrast, the residual component — variance risk that cannot be spanned by

stock returns — earns a significant premium. This suggests that options markets compensate for variance risks not present in the stock market, consistent with intermediary-based theories of option pricing (Garleanu et al. (2008), Chen et al. (2019)), and/or with the inability to hedge dynamically volatility risk during market closures.

We then reconcile our findings with the existing cross-sectional evidence. Reproducing the Ang et al. (2006) sort on our data with monthly volatility betas, we recover their negative spread on high-minus-low volatility-beta portfolios, with significant alphas relative to static factor models. Two features of their approach drive this result. First, the assumption that volatility betas are constant over a month is rejected by the data; Liao and Todorov (2026) show that the projection of volatility shocks on systematic risks varies meaningfully within a month. When we repeat the sort with daily or weekly rebalancing — frequencies at which the dynamic structure can be better accommodated — the spread in returns of volatility sorted portfolios vanishes. Second, the use of a static as opposed to a dynamic factor model when computing the alphas of the volatility sorted portfolios can also account for their magnitude. Both forces work in the same direction, and together they explain the gap between our findings and the prior literature. We further show that when expected volatility contains a nontrivial component unspanned by stock returns, the estimated volatility betas suffer from classical errors-in-variables attenuation. This attenuation inflates the apparent cross-sectional price of volatility risk whenever the volatility-mimicking portfolio fails to account for the non-spanned component properly.

The rest of the paper is organized as follows. Section 2 introduces our notation and distinguishes realized from expected variance risk. Section 3 constructs the forward-variance options portfolio. Section 4 develops the three-way decomposition, both in a general setting and within an equilibrium model. Section 5 presents our main empirical results. Section 7 contains robustness checks. Section 6 reconciles our findings with the earlier cross-sectional evidence. Section 8 concludes. Theoretical derivations and implementation details are collected in Section 9.

2 Market Variance Risks

2.1 Setting and Notation

Our interest in the paper is the volatility of the market portfolio. We denote with X_t its discounted, using the risk-free interest rate, value at time t. That is, suppose for simplicity of notation that the continuously-compounded risk-free interest rate is deterministic and its value at time t is equal to $r_{f,t}$. Then, if the value of the market portfolio at time t is $\widetilde{X}_t$, we work directly with $X_t = \widetilde{X}_t e^{-\int_0^t r_{f,u} du}$. The discounted market portfolio has the following dynamics

$$\frac{dX_t}{X_{t-}} = m_t dt + \sigma_t^\top dW_t + \int_E (e^{\gamma_x(z)} - 1)(\mu - \nu^{\mathbb{P}})(dt, dz), \tag{1}$$

where W_t is K-dimensional Brownian motion that captures systematic diffusive risks in the market portfolio, for some positive integer K, and μ is an integer-valued measure that counts the jumps in $\log(X)$. Since we will consider multiple assets, the jump measure is defined on some (large) space E and $\gamma_z : E \to \mathbb{R}$ maps the jumps in that space to jumps in X. The jump compensator is denoted with $\nu^{\mathbb{P}}$ and takes the form $\nu^{\mathbb{P}}(dt, dz) = dt \otimes \nu_t^{\mathbb{P}}(dz)$, for some measure $\nu_t^{\mathbb{P}}$ on E that captures the jump intensity at time t. Finally, m_t and σ_t are some processes that capture the instantaneous drift (mean) and volatility of X, respectively.

The above specification of the market portfolio dynamics is very general and covers most reduced-form or equilibrium-based models considered in prior work. The use of higher-dimensional Brownian motion in the modeling allows us to consider different potential sources of systematic risk in the market portfolio that require different compensation. For example, in long-run risk models, these can include separate shocks to aggregate consumption growth, its instantaneous mean and volatility, see e.g., Eraker and Shaliastovich (2008) and Eraker and Yang (2022). Similarly, in the framework of Campbell and Vuolteenaho (2004), the risk of the market portfolio can be related to separate systematic shocks in the form of cash flow and discount rate news.

The aggregate pricing kernel in the economy is denoted by M_t and has the following dynamics

$$\frac{dM_t}{M_{t-}} = \lambda_t^\top dW_t + \int_E (Y - 1)\mu(dt, dz), \tag{2}$$

where λ_t is a $K \times 1$ vector capturing the prices of risk for W_t and $Y : \mathbb{R} \times E \to \mathbb{R}$ is a

function capturing the prices of risk for the jumps. The pricing kernel induces risk-neutral probability measure, denoted with $\mathbb{Q}$. The dynamics of X_t under $\mathbb{Q}$ is given by

$$\frac{dX_t}{X_{t-}} = \sigma_t^\top dW_t^{\mathbb{Q}} + \int_E (e^{\gamma_x(z)} - 1)(\mu - \nu^{\mathbb{Q}})(dt, dz), \tag{3}$$

where $W_t^{\mathbb{Q}}$ is a $\mathbb{Q}$-Brownian motion and $\nu^{\mathbb{Q}}(dt, dz) = Y(t, z)\nu^{\mathbb{P}}(dt, dz) = dt \otimes \nu_t^{\mathbb{Q}}(dz)$ is the jump compensator under $\mathbb{Q}$.

2.2 Expected versus Realized Variance Risk

We proceed with introducing our notions of variance risk whose pricing we will study. We start with our measure of realized variance risk. We consider the following one

$$\begin{aligned} RV_{t,T} &= 2\,\mathrm{plim}_{\Delta\downarrow 0} \sum_{i=1}^{\lfloor T/\Delta \rfloor} \left(\frac{X_{t+i\Delta} - X_{t+(i-1)\Delta}}{X_{t+(i-1)\Delta}} - \log\left(\frac{X_{t+i\Delta}}{X_{t+(i-1)\Delta}} \right) \right) \\ &= \int_t^{t+T} \sigma_s^\top \sigma_s ds + 2 \int_t^{t+T} \int_E (e^{\gamma_x(z)} - \gamma_x(z) - 1)\mu(ds, dz). \end{aligned} \tag{4}$$

Each summand in the first line of the above display is the difference between simple and log-returns.[3] The above notion of realized variance is slightly different from the usual one which involves summation of squared high-frequency log-returns, see e.g., Andersen et al. (2003). The only difference of $RV_{t,T}$ above from the standard definition of realized variance is due to the jumps, as clear from the following expansion:

$$2\left(\frac{X_{t+\Delta} - X_t}{X_t} - (x_{t+\Delta} - x_t) \right) = (x_{t+\Delta} - x_t)^2 + O((x_{t+\Delta} - x_t)^3), \text{ as } \frac{X_{t+\Delta}}{X_t} \to 1, \tag{5}$$

where we used the short-hand notation $x_t = \log(X_t)$. The compensation for the realized variance risk is referred to in the literature as variance risk premium, see e.g., Bollerslev et al. (2009), Carr and Wu (2009) and Drechsler and Yaron (2011). We will refer to it as *realized variance risk premium* to highlight the fact that it is compensation for exposure to *realized variance risk* as opposed to *expected variance risk* that we consider below. Formally,

[3]The difference between simple and log-returns has been used as a way to test for presence of jumps by Jiang and Oomen (2008).

we define the realized variance risk premium as:

$$\mathrm{RVP}_{t,T} = \frac{1}{T}(\mathbb{E}_t^{\mathbb{P}} - \mathbb{E}_t^{\mathbb{Q}})(RV_{t,T}). \tag{6}$$

This risk premium can be related (approximately) to the expected payoff of a long position in a variance swap. An alternative measure of variance is the instantaneous expected variance that we define as

$$\begin{aligned} V_t &= \lim_{T \downarrow 0} \frac{1}{T} \mathbb{E}_t^{\mathbb{Q}} \left(RV_{t,T} \right), \text{ a.s.}, \\ &= \sigma_t^\top \sigma_t + 2 \int_E (e^{\gamma_x(z)} - \gamma_x(z) - 1) \nu_t^{\mathbb{Q}}(dz). \end{aligned} \tag{7}$$

V_t is the spot counterpart of the (squared) VIX volatility index which equals $\frac{1}{T}\mathbb{E}_t^{\mathbb{Q}}\left(RV_{t,T}\right)$, for T corresponding to one month. Unlike realized variance, V_t is a measure of *conditional expected variance* and constitutes a state variable that captures changes in the investment opportunity set that are of direct relevance for the ICAPM model and also for most dynamic equilibrium asset pricing models.

We note that shocks to realized variance and to expected variance differ in general. To see this, note that we can write the following for the realized variance risk premium:

$$\begin{aligned} \mathrm{RVP}_{t,T} = &\frac{1}{T} \int_t^{t+T} (\mathbb{E}_t^{\mathbb{P}} - \mathbb{E}_t^{\mathbb{Q}})(V_s - V_t) ds \\ &+ \frac{2}{T} \mathbb{E}_t^{\mathbb{P}} \left(\int_t^{t+T} \int_E (e^{\gamma_x(z)} - \gamma_x(z) - 1)(\nu_s^{\mathbb{P}}(dz) - \nu_s^{\mathbb{Q}}(dz)) ds \right). \end{aligned} \tag{8}$$

The first component of the realized variance risk premium is due to compensation for shocks to instantaneous expected variance. The second component is due to jump risk and will be present even if the investment opportunity set was static, e.g., in an economy with i.i.d. market returns. In fact, the jump part of $\mathrm{RVP}_{t,T}$ dominates for T small as is clear from the following:

$$\mathrm{RVP}_t = \lim_{T \downarrow 0} \mathrm{RVP}_{t,T} = 2 \int_E (e^{\gamma_x(z)} - \gamma_x(z) - 1)(\nu_t^{\mathbb{P}}(dz) - \nu_t^{\mathbb{Q}}(dz)). \tag{9}$$

Our interest in this paper will be in the compensation for shocks to the instantaneous ex-

pected variance V_t, which we refer to as *expected variance risk premium*. It is formally defined as

$$\mathrm{EVP}_t = \frac{1}{T} \lim_{T \downarrow 0} [\mathbb{E}_t^{\mathbb{P}}(V_{t+T} - V_t) - \mathbb{E}_t^{\mathbb{Q}}(V_{t+T} - V_t)]. \tag{10}$$

EVP_t can be readily linked with the prices of risk entering the dynamics of the market portfolio X, i.e., with W and μ. To do this, we need some more notation. In particular, the dynamics of X under $\mathbb{P}$ and $\mathbb{Q}$ imply the following dynamics for V_t under $\mathbb{P}$:

$$dV_t = \alpha_t^{\mathbb{P}} dt + \eta_t^\top dW_t + \int_E \gamma_v(z)(\mu - \nu^{\mathbb{P}})(dt, dz), \tag{11}$$

for some coefficients $\alpha_t^{\mathbb{P}}$ and η_t as well as a function $\gamma_v : \mathbb{R}_+ \times E \to \mathbb{R}$. We have similarly for the dynamics of V_t under the risk-neutral measure:

$$dV_t = \alpha_t^{\mathbb{Q}} dt + \eta_t^\top dW_t^{\mathbb{Q}} + \int_E \gamma_v(z)(\mu - \nu^{\mathbb{Q}})(dt, dz), \tag{12}$$

for some coefficient $\alpha_t^{\mathbb{Q}}$. The coefficients in the dynamics of V_t can be easily expressed as transforms of those in the dynamics of X_t but we will not do this here in order to simplify exposition. The difference in the drift coefficients of V_t under $\mathbb{P}$ and $\mathbb{Q}$ is the instantaneous risk premium for it:

$$\mathrm{EVP}_t = \alpha_t^{\mathbb{P}} - \alpha_t^{\mathbb{Q}} = \eta_t^\top \lambda_t + \int_E \gamma_v(z)(\nu^{\mathbb{P}} - \nu^{\mathbb{Q}})(dz). \tag{13}$$

In words, EVP_t reflects compensation for diffusive and jump risks in the instantaneous expected variance. Recall in this regard that $\gamma_v(z)$ captures the jumps in V and not those in the market index.

3 Spanning Expected Variance via Forward-Variance Options Portfolio

Our goal next is to link the realized and expected variance risks to the risks of holding portfolios of short-dated European-style options. This will allow us to study the risk premium associated with these risks using directly the excess returns of their replicating option port-

folios. We start with constructing a portfolio that spans realized variance risk. This portfolio is simply the one behind the construction of the VIX volatility index, see e.g., Britten-Jones and Neuberger (2000) and Carr and Wu (2009). More specifically, let us denote with $C_{s,t}(K)$ and $P_{s,t}(K)$ the prices of call and put options with strike K written on the asset at time s that expire at some future date $t > s$. Denote further with $F_{s,t}$ the value of a future contract written on the underlying asset at time s that expires at time t in the future. Finally, let $R_{f,s,t}$ be the risk-free factor over the period $[s, t]$, given by

$$R_{f,s,t} = e^{\int_s^t r_{f,u} du}. \tag{14}$$

One dollar invested in a risk-free asset at time s grows to $R_{f,s,t}$ at time t. With this notation, we then have (see Bakshi and Madan (2000) and Carr and Madan (2001)):

$$V_{s,T} = \mathbb{E}_s^{\mathbb{Q}} \left(\int_s^T V_u du \right) = 2R_{f,s,T} \int_0^{F_{s,T}} \frac{P_{s,T}(K)}{K^2} dK + 2R_{f,s,T} \int_{F_{s,T}}^{\infty} \frac{C_{s,T}(K)}{K^2} dK. \tag{15}$$

The right-hand-side of the above expression is simply the value at time s of a portfolio of out-of-the-money options over a continuum of strikes, with amount invested in the option with strike K being equal to $2R_{f,s,T}/K^2$. Denote the excess return of this portfolio over $[s, t]$, for $s < t \leq T$, with $R^{rv}_{s,t}$. If we hold the portfolio until expiration, then its excess return over the risk-free rate, is given by

$$R^{rv}_{s,T} = \frac{2R_{f,s,T}}{V_{s,T}} \left[\frac{F_{T,T} - F_{s,T}}{F_{s,T}} - \log \left(\frac{F_{T,T}}{F_{s,T}} \right) \right] - R_{f,s,T}, \tag{16}$$

and we can show that

$$\mathbb{E}_s^{\mathbb{P}}(R^{rv}_{s,T}) = \mathrm{RVP}_s/V_s + O_P(T - s), \text{ as } T \downarrow s. \tag{17}$$

The expansion in (17) implies that the expected excess return per unit time is approximately RVP_s/V_s, which by equation (9) is pure price jump compensation. More generally, as shown in the proof of Proposition 2 below, we have

$$R^{\mathrm{rv}}_{s,t} = \frac{R_{f,s,t}}{V_{s,T}} \left(V_{t,T} - V_{s,T} + 2 \left[\frac{F_{t,T}}{F_{s,T}} - 1 - \log \left(\frac{F_{t,T}}{F_{s,T}} \right) \right] \right). \tag{18}$$

This return depends both on price shocks and expected variance shocks. The former dominate the latter when T goes to zero and we have the analogue of (17), see Proposition 2 below.

We continue next with linking shocks to expected instantaneous variance V_t with the risk of a portfolio of available European-style short-dated options. This is significantly more difficult to do than the construction of a portfolio of options that spans realized variance risk. Essentially, we need a way to hedge the price jump risk that the above portfolio of options is exposed to and we want to do this in a fully model-free way. To achieve this goal, we will construct a portfolio of options that spans short-term forward variance. Towards this end, we first denote the following risk-neutral expectation of future variance:

$$V_{s,T,2T} = \mathbb{E}_s^{\mathbb{Q}} \left(\int_T^{2T} V_u du \right), \ s \in [0, T), \ T > 0. \tag{19}$$

In our application, T will be at most one week, and in the last three years of the sample, it will be one day. Given our analysis above, we can write

$$\begin{aligned} V_{s,T,2T} = {} & 2R_{f,s,2T} \int_0^{F_{s,2T}} \frac{P_{s,2T}(K)}{K^2} dK + 2R_{f,s,2T} \int_{F_{s,2T}}^{\infty} \frac{C_{s,2T}(K)}{K^2} dK \\ & - 2R_{f,s,T} \int_0^{F_{s,T}} \frac{P_{s,T}(K)}{K^2} dK - 2R_{f,s,T} \int_{F_{s,T}}^{\infty} \frac{C_{s,T}(K)}{K^2} dK. \end{aligned} \tag{20}$$

That is, $V_{s,T,2T}$ is the value at time s of a portfolio that is long options expiring at time $2T$ and short options expiring at time T. The amount invested in the longer-dated option with strike K is $2R_{f,s,2T}/K^2$ and the amount of shorter-dated option with strike K that is shorted is $2R_{f,s,T}/K^2$. We refer to this portfolio as the *forward-variance spanning portfolio* and denote its excess return over $[s, t]$, for $s < t \leq T$, with $R^{ev}_{s,t}$.

We can show that the payoff of this option portfolio at some time $t \in (s, T)$ is equal to $R_{f,s,t} V_{t,T,2T}$, i.e., it equals the value of the same position in out-of-the-money options at time t times the risk-free adjustment factor. Note that, the options behind $V_{s,T,2T}$ and $V_{t,T,2T}$ in general differ because of the change in the underlying asset price. For example, if $F_{t,T} > F_{s,T}$, the out-of-the-money options for strikes in the range $(F_{s,T}, F_{t,T})$ in the portfolio created at time s are puts while those in the portfolio created at time t are calls. In spite of these differences in the portfolios behind $V_{s,T,2T}$ and $V_{t,T,2T}$, we show in the next proposition that

the value of the portfolio at time t that is started at time s is exactly $R_{f,s,t}V_{t,T,2T}$.

Proposition 1 (Forward-variance spanning)**.** *Let $V_{s,T,2T}$ denote the value at time s of the option portfolio defined in equation (20), which is long out-of-the-money options expiring at $2T$ with weights $2R_{f,s,2T}/K^2$ and short out-of-the-money options expiring at T with weights $2R_{f,s,T}/K^2$. Then, for any $t \in [s,T)$, the value at time t of this portfolio, held without rebalancing from s to t, equals*

$$\Pi_{s,t} = R_{f,s,t} \cdot V_{t,T,2T}. \tag{21}$$

In particular, the holding-period excess return of the portfolio over $[s,t]$ is

$$R^{ev}_{s,t} = R_{f,s,t}\frac{V_{t,T,2T}}{V_{s,T,2T}} - R_{f,s,t}. \tag{22}$$

The result in Proposition 1 holds only for the forward-variance spanning portfolio behind $V_{s,T,2T}$. A similar result does not hold for its components that span $V_{s,T}$ and $V_{s,2T}$, and this is the reason why we use the forward-variance spanning portfolio to study expected variance risk.

We note that Proposition 1 holds for any value of T and shows that the forward-variance portfolio is exposed only to expected variance risk. In particular, equation (22) shows that $R^{ev}_{s,t}$ depends on the price path only through its effect on $V_{t,T,2T}$: any price move, jump or diffusive, that leaves $V_{t,T,2T}$ unchanged leaves $R^{ev}_{s,t}$ unchanged as well, while a price move accompanied by a change in $V_{t,T,2T}$ — as would occur in models with correlated price-variance jumps — is correctly registered as variance risk. To link further the return $R^{ev}_{s,t}$ to shocks to the instantaneous expected variance V_t, we need T small. The formal result is given in the next proposition:

Proposition 2 (Expected and realized variance risk premia from option returns)**.** *Under the dynamics in equations (1)–(3) and standard integrability conditions, the holding-period expected excess returns of the realized-variance and forward-variance spanning portfolios satisfy, for any $0 \le s < t < T$ and as $T \downarrow 0$,*

$$\frac{T-s}{t-s}\mathbb{E}^{\mathbb{P}}_s(R^{rv}_{s,t}) = RVP_s/V_s + O_P(T), \text{ for } 0 \le s < t \le T \text{ and as } T \downarrow 0, \tag{23}$$

and

$$\frac{1}{t-s}\mathbb{E}_s^{\mathbb{P}}(R_{s,t}^{ev}) = EVP_s/V_s + O_P(T), \text{ for } 0 \leq s < t \leq T \text{ and as } T \downarrow 0, \tag{24}$$

where RVP_s is the instantaneous realized variance risk premium defined in equation (9) and EVP_s is the instantaneous expected variance risk premium defined in equation (10).

To illustrate the difference between shocks to realized and expected variances, and between the returns of the corresponding options portfolios that replicate these risks, we compute $R_{s,t}^{rv}$ and $R_{s,t}^{ev}$ for different scenarios for price and underlying (latent) variance changes in a standard one-factor Heston volatility model plus price jumps, with jump intensity proportional to the diffusive spot variance. We do not specify the distribution of the price jumps as this is immaterial for the illustration here. In this model, there are two state variables driving the dynamics of the market portfolio: the spot price and the spot diffusive variance.

We consider moves in price and variance which are one and four times the typical downside and upside moves in the market index and the variance, respectively, observed in the real data used in our application. We compute within the model the values of $V_{t,T}$, $V_{t,2T}$ and $V_{t,T,2T}$ for a value of T of one trading day, and from there $R_{s,t}^{rv}$ and $R_{s,t}^{ev}$ using (18) and (22). Note that for these calculations, we only need the size of the moves in X and V as well as the model for the dynamics of V under $\mathbb{Q}$.

The results are reported in Table 1. Consistent with our analysis above, a significant part of realized variance risk is actually price (jump) risk. On the other hand, the forward-variance spanning portfolio is immune to price risk that does not carry information about expected variance, and only reacts to shocks to V_t, which we are after. This holds regardless of how large the price move is or whether it originates from the diffusive or jump component of the price process, since $R_{s,t}^{ev}$ depends on the price path only through $V_{t,T,2T}$ (equation (22)). Note that there is a slight difference between $V_t/V_s - 1$ and $R_{s,t}^{ev}$. This is due to the use of a finite T (Proposition 2 is an asymptotic result for shrinking T). Nevertheless, the two quantities are highly correlated. Overall, this table shows that the forward-variance options portfolio decouples price and variance shocks.

Table 1: Sensitivity of Variance Portfolios to Price and Variance Changes

Price Change (%) $(X_t/X_s - 1)$	**Variance Change (%)** $(V_t/V_s - 1)$	$R^{rv}_{s,t}(\%)$	$R^{ev}_{s,t}(\%)$
-0.06%	0.0%	0.82%	0.00%
-0.25%	0.0%	6.61%	0.00%
0.00%	+1.6%	0.28%	1.53%
0.00%	+6.4%	4.93%	6.11%
-0.06%	+1.6%	0.73%	1.53%
-0.25%	+6.4%	12.83%	6.11%

Note: $s = 0$, $t - s$ corresponds to 5-minute interval, one trading day has 6.5 hours, T equals one trading day; $V_0 = 0.02$ and $r_{f,t} = 0$. The numbers are computed for a square-root diffusion model for V_t with mean-reversion parameter equal to 8 (corresponding to a half-life of a shock of roughly one month) and mean value of 0.02 (reported in annualized units).

To conclude, the forward-variance spanning portfolio constructed above provides a fully model-free, tradable proxy for shocks to expected market variance, and as such addresses the central limitations of the three approaches discussed in the introduction. Unlike parametric option-pricing studies, the construction requires no assumption on the dynamics of V_t or on the form of the pricing kernel; the spanning result in equation (21) holds under the general specification in equations (1)-(3). Unlike delta-hedged option gains, the portfolio's return is exposed only to shocks to expected variance and is immune to price jumps as shown in Proposition 1 and illustrated in Table 1. And unlike realized-variance-based measures, whose risk premium reduces to pure jump compensation as the measurement horizon shrinks (equation (9)), our portfolio isolates expected-variance risk that is of direct relevance for the ICAPM.

4 Decomposing Expected Variance Risk

To understand better what risks in V_t are priced, we can combine the position in the forward-variance options portfolio with a dynamic position that involves either holding the underlying asset (the market portfolio) or a wide cross-section of stocks. We first conduct the analysis in a general model-free setting and then provide details about its feasible implementation

using real data. We finish with an illustration of the decomposition of risks to V_t within a standard dynamic asset pricing equilibrium model.

4.1 Model-free Decomposition

We denote the prices of a cross section of N stocks at time t, discounted at the risk-free interest rate, with the $N \times 1$ vector Y_t. Then, the cumulative gains of our portfolio of options, held over the period $[s,t]$, and a dynamic portfolio in the cross-section of stocks with weights at time $u \in [s,t]$ given by the vector ϕ_u is given by

$$R_{f,s,t}V_{t,T,2T} - V_{s,T,2T} - \int_s^t \phi_u^\top (dY_u + r_{f,u}du). \tag{25}$$

Different choices for ϕ will allow us to relate the risks in V_t to the systematic risks in the individual stocks and those in the market portfolio. To show this, we introduce the dynamics of the individual asset prices. The j-th element of Y_t evolves according to

$$\frac{dY_t^{(j)}}{Y_{t-}^{(j)}} = m_t^{(j)}dt + \beta_t^{(j)\top}dW_t + d\widetilde{W}_t^{(j)} + \int_E (e^{\gamma_j(z)} - 1)(\mu - \nu^{\mathbb{P}})(dt, dz),\ j = 1, ..., N, \tag{26}$$

where W_t is the vector of Brownian motions that appears in the dynamics of the market portfolio and $\widetilde{W}_t^{(j)}$ are independent Brownian motions capturing idiosyncratic diffusive risks in the individual assets. We note that our setup is quite general. In particular, we allow for the situation in which the set of systematic risk factors that drive the cross-section of individual stock returns and the ones driving the instantaneous expected market variance to differ.

We define with ϕ_u^* the position in the underlying stocks at time u that minimizes the diffusive risk in the forward-variance spanning portfolio. That is,

$$\phi_\cdot^* = \underset{\phi}{\text{argmin}} \left[R_{f,s,\cdot}V_{\cdot,T,2T} - V_{s,T,2T} - \int_s^\cdot \phi_z^\top dY_z,\ R_{f,s,\cdot}V_{\cdot,T,2T} - V_{s,T,2T} - \int_s^\cdot \phi_z^\top dY_z \right]^c, \tag{27}$$

where minimization is over cadlag predictable processes ϕ and where $[X,Y]^c$ denotes the continuous part of the quadratic variation between two generic processes X and Y. This simply means that ϕ_u^* minimizes diffusive instantaneous variance of the hedged portfolio at

any point u. We can equivalently define ϕ_u^* as

$$\left[R_{f,s,u}V_{u,T,2T} - V_{s,T,2T} - \int_s^u \phi_z^{*\top} dY_z, \int_s^u \phi_z^{*\top} dY_z\right]^c = 0, \quad s \leq u \leq t, \tag{28}$$

In other words, ϕ_u^* is the local projection coefficient of the diffusive variance increment onto the full cross-section of stock returns. In a finite cross-section, this projection would load on both the systematic and the idiosyncratic components of individual stock returns. We work in a large-N setup, however, in which the idiosyncratic components in the cross-section average out, and the projection in equation (28) is therefore a projection onto the systematic factor space. As we will see in the next section, ϕ_u^* can be easily estimated from the data.

We can view $\int_s^u \phi_z^{*\top} dY_z$ as the capital gains at time u of a market variance-mimicking portfolio constructed from the cross-section of available stocks. We denote the cumulative excess return of a portfolio involving this dynamic strategy in the individual stocks plus a money-market account as

$$R_{s,t}^{cs} = \frac{\int_s^t \phi_u^{*\top} dY_u}{V_{s,T,2T}}. \tag{29}$$

We use the money-market account to make the initial investment in the portfolio equal exactly to $V_{s,T,2T}$. Similarly, we denote the cumulative excess return of a portfolio that is exposed to the residual market variance risk as

$$R_{s,t}^{\perp cs} = \frac{R_{f,s,t}V_{t,T,2T} - \int_s^t \phi_u^{*\top} dY_u}{V_{s,T,2T}} - R_{f,s,t}. \tag{30}$$

We can further decompose the variance-mimicking portfolio into a dynamic portfolio in the market index, that provides a best hedge for it, and one orthogonal to it. Towards this end, let φ_u^* be the position in the market portfolio at time u such that

$$\left[\int_s^u \phi_z^{*\top} dY_z - \int_s^u \varphi_z^* dX_z, \int_s^u \varphi_z^* dX_z\right]^c = 0, \quad s \leq u < T. \tag{31}$$

Then, we denote

$$R_{s,t}^m = \frac{\int_s^t \varphi_u^* dX_u}{V_{s,T,2T}} \text{ and } R_{s,t}^{\perp m} = \frac{\int_s^t \phi_u^{*\top} dY_u - \int_s^t \varphi_u^* dX_u}{V_{s,T,2T}}. \tag{32}$$

Altogether, we have the following decomposition

$$R^{ev}_{s,t} = R^{m}_{s,t} + R^{\perp m}_{s,t} + R^{\perp cs}_{s,t}. \tag{33}$$

The corresponding decomposition of the change in variance is given by

$$R_{f,s,t}V_{t,T,2T} - V_{s,T,2T} = \Delta V^{m}_{s,t} + \Delta V^{\perp m}_{s,t} + \Delta V^{\perp cs}_{s,t}. \tag{34}$$

The returns on the right-hand side of (33) are those of dynamic portfolios that involve the short-dated options, the market index and the cross-section of stocks, provided one knows ϕ^*_u and φ^*_u. We will show in the next section how one can form estimates of the latter. The excess returns of the three portfolios can tell us the sources of risk in expected market variance that are priced by investors and how they relate to other risks of importance, i.e., the ones priced in the cross-section of stock returns and the market portfolio. In particular, $R^{\perp m}_{s,t}$ and $R^{\perp cs}_{s,t}$ contain risks that cannot be dynamically spanned/hedged by investing in the market portfolio or the entire cross-section of stocks, respectively.

It is important to note that we allow for ϕ^*_u and φ^*_u to vary over the interval $[s, t]$. This will capture, in particular, the situation in which the exposure of V_t to the latent systematic risks embedded in the individual stock prices can vary over time. Such a situation will naturally occur when V_t is a non-linear function of systematic shocks that drive the cross-section of stock returns, as is the case for most equilibrium models. In such a situation, if we were to restrict ϕ^*_u to be constant, then $R^{\perp cs}_{s,t}$ will be non-zero even though the risks in V_t can be replicated by a dynamic portfolio of stocks.

The decomposition in (33) is based on the diffusive risks in V_t and in the individual stock returns. The reason for not including jumps in the analysis is that a model-free decomposition of jump risks, like the one for the diffusive risks, is not possible. This is pretty intuitive as jumps are rare events and they might not even occur over a short period of time. Moreover, the relation between jumps in the assets and V_t can differ, depending on the size of the jumps. That said, it is natural to expect that the same decomposition as for the diffusive risks above will also hold at least for the small jumps. If this is not the case, then $R^{\perp cs}_{s,t}$ can include jumps that arrive at the same time with the jumps in the systematic risk factors driving the cross-section of returns. These jumps in $R^{\perp cs}_{s,t}$ will reflect a residual component due to the nonlinear relation between the jumps in V_t and those in the stock

prices. Nevertheless, the reason for the existence of $R^{\perp cs}_{s,t}$ remains unchanged: it reflects the existence of shocks to expected variance that cannot be dynamically hedged via portfolios of stocks.

4.2 Feasible Implementation

Our goal next is to make the decomposition in (34) feasible. We will do this using high-frequency observations of stocks and index options. We provide the main ideas in this section and then we cover the implementation details in the Appendix. Our inference for the projection coefficients (portfolio weights) ϕ^*_t and φ^*_t will be based on high-frequency observations at intervals of length Δ_n, for some small $\Delta_n > 0$ (and much smaller than T), of $V_{t,T}$, X_t and Y_t.

We start with ϕ^*_t. We will estimate this quantity using blocks of high-frequency increments of $V_{t,T}$ and Y_t. We will pretend that a counterpart of ϕ^*_t corresponding to log-transforms of the variables remains constant over the estimation window. In this case, ϕ^*_t can be simply determined by the vector of projection coefficients of log-variance increments on high-frequency log-returns of systematic risk factors uncovered from the cross-section of stocks.

Since we do not assume prior knowledge of the factors driving the cross-section of stock returns, we estimate these via Principal Component Analysis (PCA). The details are provided in the Appendix. The PCA recovers the latent factors up to a rotation but this indeterminacy is irrelevant for the analysis here as we are only interested in the projection of our variance increments on the latent factors, see e.g., Giglio and Xiu (2021). We note that the recovered systematic risk factors are simply portfolios of stock returns.

Finally, the estimate of φ^*_t can be obtained similarly by running a linear regression of the projection of the log-variance increments on the systematic factors obtained above on the log-increments of the market portfolio. To avoid biases in our expected return estimates due to the estimation error in recovering ϕ^*_t and φ^*_t, the returns $R^m_{s,t}$, $R^{\perp m}_{s,t}$ and $R^{\perp cs}_{s,t}$ are computed using estimates of ϕ^*_t and φ^*_t from a time window prior to time s. The details for all computations are provided in the Appendix.

4.3 Decompositions in an Equilibrium Model Example

We can illustrate the decomposition of $R^{ev}_{s,t}$ within an equilibrium model with a representative agent. We will use the model in Eraker and Yang (2022), which was developed with the goal of explaining volatility-related products such as options on the VIX index, and is therefore particularly well suited for our analysis here. In the model, there is a representative agent endowed with Duffie-Epstein utility preferences. The aggregate consumption of the agent evolves according to

$$d\log(C_t) = \mu_c dt + \sigma^C_t dB^C_t, \tag{35}$$

where C_t is aggregate consumption and B^C_t is a Brownian motion capturing shocks to consumption growth. The variance of consumption growth follows

$$d(\sigma^C_t)^2 = \kappa_v(\theta_v - (\sigma^C_t)^2)dt + \sigma_v\sigma^C_t dB^v_t + \zeta_v dN_t, \tag{36}$$

where B^v_t is a Brownian motion independent of B^C_t that captures diffusive shocks to volatility of consumption growth, ζ_v is a time-invariant random variable that captures jumps in consumption growth volatility, N_t is a compound Poisson process that counts the jumps in σ^C_t, with intensity parameter λ_t evolving according to

$$d\lambda_t = \kappa_\lambda(\theta_\lambda - \lambda_t)dt + \sigma_\lambda\sqrt{\lambda_t}dB^\lambda_t, \tag{37}$$

where B^λ_t is a Brownian motion independent of B^C_t and B^v_t.

In this model, the vector W_t driving the aggregate pricing kernel in (2) is given by $(B^C_t, B^v_t, B^\lambda_t)$ while μ corresponds to the counting measure for the jumps in the consumption volatility. In other words, shocks to aggregate consumption, consumption volatility and intensity of jumps in consumption volatility are the priced risks in this economy.

We model the dividend of asset j as

$$D^{(j)}_t = S^{(j)}_t \times C_t, \tag{38}$$

where $S_t = \left(S^{(1)}_t, ..., S^{(N}_t\right)^\top$ is an N-dimensional vector of shares of individual stocks' dividends in aggregate consumption, with $\sum_{j=1}^N S^{(j)}_t = 1$. This assumes that aggregate consumption equals aggregate dividends. We can also introduce a wedge between aggregate

consumption and aggregate dividends, as done in prior work. For simplicity, we do not do this here, as such an extension is not important for the purposes of our analysis.

Given the pricing kernel in the model, the specification for the aggregate consumption and the dividends of individual stocks, we have

$$Y_t^{(j)} = G_j(S_t, C_t, \sigma_t^C, \lambda_t), \tag{39}$$

for some function G_j, provided the vector $(S_t, C_t, \sigma_t^C, \lambda_t)$ is a Markov process. Similarly, we can write for the market portfolio

$$X_t = C_t \times G(\sigma_t^C, \lambda_t), \tag{40}$$

for some function G.

The instantaneous market variance (V_t) in this model is some function of (σ_t^C, λ_t), i.e., it is a function of the two state variables that drive the consumption growth dynamics. As a result, compensation for shocks to these two state variables generates the expected market variance risk premium (EVP_t). The return on the market portfolio, on the other hand, is exposed to shocks to $(\log(C_t), \sigma_t^C, \lambda_t)$. Therefore, the risks in market returns and market variance overlap only partially. The individual stocks are exposed to shocks to $(\log(C_t), \sigma_t^C, \lambda_t)$ but also to the state vector of dividend shares S_t. If there is a difference in the individual stocks' exposure to σ_t^C and λ_t in the cross-section, then $R_{s,t}^{\perp cs}$ should contain no diffusive risk. In other words, we should be able to span the diffusive risks in the market variance by dynamic portfolios of individual stocks. If stocks' exposures to σ_t^C and λ_t are perfectly correlated, on the other hand, then there will be risks in the market variance that cannot be spanned by trading dynamically in individual stocks, and therefore $R_{s,t}^{\perp cs}$ will be non-negligible.

Overall, this example shows that even in a relatively simple representative agent model, the risks in the market portfolio, its variance and the cross-section of stocks can differ. Naturally, the wedge between market variance risk and systematic risks embedded in the cross-section of assets will grow larger in economies with heterogeneous agents, and in particular in settings in which the marginal investor in the options market can differ from the one in the stock market, see e.g., Garleanu et al. (2008), Chen et al. (2019), Almeida and Freire (2026) and Ober (2024). In such settings, one would expect $R_{s,t}^{\perp cs}$ to be large and

non-trivial. We will evaluate this empirically later.

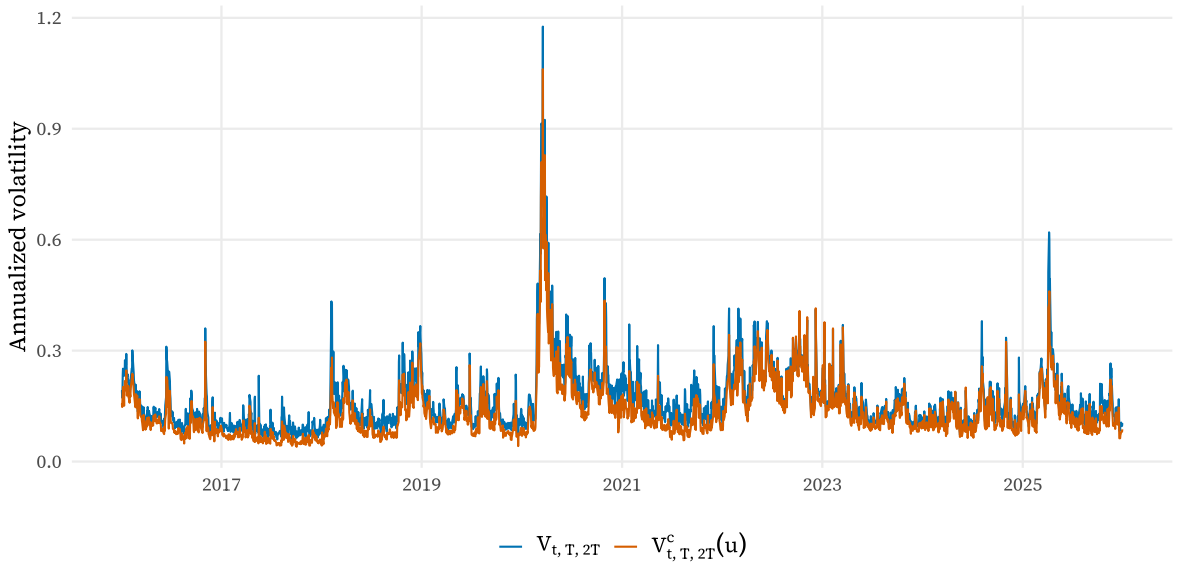
5 Empirical Evidence for the Pricing of Expected Market Variance Risk

5.1 Data

We collect data from several sources for our empirical study. First, we use high-frequency S&P 500 index option data from CBOE Datashop. The sample period for the option data is January 4, 2016 until December 31, 2025. For each trading day, we sample option prices at five-minute frequency over the interval 9.35-16.00 EST. This results in 78 intraday variance increments. We also collect an overnight variance increment over the period from 16.00 EST on a given trading day until 9.35 EST on the following trading day. On each day, we look for options with the shortest available tenors but such that there are at least 2 out-of-the-money puts and calls at each point in time during the day. The minimum tenor at the end of each interval should be at least one trading day. The moneyness for each maturity and each point in time is determined by a synthetic forward extracted from near-the-money options by an application of put-call parity. This way of selecting the tenors used in the calculations leads to sample medians of 2 and 4 business days for the first and second tenor, respectively, at the end of each trading day. Finally, the risk-free interest rate used for computing $R_{f,s,t}$ is the 3-month Treasury Bill rate obtained from the FRED database. The calculation of $V_{s,t}$ is done by a Riemann sum approximation of the integrals in equation (15) from the available option mid-quotes (as is standard practice we consider only strikes with positive bids for the out-of-the-money options).

In Figure 1, we plot the estimated forward variance $V_{t,T,2T}$ (and an alternative variance proxy that we will consider in our robustness checks section) at the end of each trading day. As seen from the figure, our sample covers periods with very high and very low volatility. In particular, our sample covers the Covid pandemic and the recent period of trade war tensions, both of which have generated turbulence on the stock markets and high volatility. Figure 1 also reveals both high and low frequency variation in market volatility.

Figure 1: **Measures of Expected Market Volatility**



The figure plots $V_{t,T,2T}$ (blue line) and $V^c_{t,T,2T}(u)$ (red line) at the end of each day in the sample. The series are reported in annualized volatility units.

The second data set used in our study is high-frequency data for the 500 largest stocks traded in the US and for the VTI ETF that tracks the CRSP total market index (which is our proxy for the market index portfolio). The source of the data is the TAQ database. The cross-section of stocks changes on an annual basis, based on the stocks' market capitalization at the end of the previous year.

5.2 Expected Variance Risk Premium and Its Decomposition

Measurement error in options is nontrivial and, as recently pointed out by Duarte et al. (2024), can have a nontrivial effect on inference for the variance risk premium. In our case, noise in option prices can have an inverse effect in the form of bias in the calculation of the returns as we need to divide by a noisy measure when computing the returns on the options portfolio. To mitigate this effect, we work with the following aggregated returns:

$$\overline{R}^{ev}_{s,t} = \frac{\sum_i [R_{f,s_i,t_i}(V_{t_i,T_i,2T_i} - V_{s_i,T_i,2T_i})]}{\sum_i V_{s_i,T_i,2T_i}}, \tag{41}$$

where the summation above is over all subintervals $[s_i, t_i]$ that are in $[s, t]$. We denote similarly $\overline{R}^{cs}_{s,t}$, $\overline{R}^{\perp cs}_{s,t}$ and $\overline{R}^{\perp \mathrm{m}}_{s,t}$. We compute these aggregate returns over periods of length one day when using intraday data and over five days when using overnight returns. This aggregation mitigates the impact of measurement error on our results as the division in (41) is by an average quantity.[4]

We report our estimation results for the expected variance risk premium and its decomposition in Table 2. We treat intraday and overnight returns separately in our analysis in order to allow for the possibility that the systematic risks in stocks and their connection to market variance risk can differ between periods with and without trading of stocks.[5] We report results where the number of systematic risk factors is estimated each estimation period (of length 22 days) according to a standard information criteria, see Bai and Ng (2002), as well as ones in which the number of factors is fixed to a given number.

The reported results provide strong evidence for the pricing of expected market variance risk, with the estimate for $\mathbb{E}(\overline{R}^{ev}_{s,t})$ being highly statistically significant both during intraday and overnight periods. The magnitude of the estimated $\mathbb{E}(\overline{R}^{ev}_{s,t})$ is roughly comparable with earlier model-based estimates of this quantity from options and returns data, see e.g., Table 15.1 in Singleton (2006).[6] We do not compare this number with an estimate for the realized variance risk premium since, as we discussed earlier, realized and expected variance risks differ substantially for short tenors.

[4] An alternative way, following Duarte et al. (2024), will be to bias-correct the raw returns.

[5] Bollerslev et al. (2016) and Hendershott et al. (2020) provide evidence that standard risk factors work differently during intraday and overnight periods in pricing the cross-section of stocks.

[6] When V_t follows a square-root diffusion, which is the case for many parametric models considered in prior work, then EVP_t/V_t is a constant that equals the difference in the mean-reversion coefficients of V_t under $\mathbb{P}$ and $\mathbb{Q}$.

Table 2: EVP Estimates

	Number of factors K									
	3		5		10		$K^\star$		$2K^\star$	
Panel A: Intraday										
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-4.878	(-4.53)	-4.878	(-4.53)	-4.878	(-4.53)	-4.878	(-4.53)	-4.878	(-4.53)
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.465	(-4.56)	-3.558	(-4.45)	-4.254	(-4.92)	-3.917	(-4.57)	-4.680	(-4.68)
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-3.625	(-5.59)	-3.659	(-5.56)	-3.740	(-5.59)	-3.751	(-5.58)	-3.791	(-5.57)
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	0.159	(0.55)	0.100	(0.26)	-0.514	(-1.08)	-0.166	(-0.35)	-0.889	(-1.19)
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-1.413	(-1.46)	-1.320	(-1.32)	-0.624	(-0.60)	-0.961	(-0.94)	-0.198	(-0.18)
Panel B: Overnight										
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-5.573	(-4.93)	-5.573	(-4.93)	-5.573	(-4.93)	-5.573	(-4.93)		
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-4.676	(-4.17)	-4.687	(-4.03)	-2.772	(-3.16)	-3.334	(-4.28)		
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-4.006	(-4.10)	-4.139	(-4.17)	-4.273	(-4.25)	-4.115	(-4.13)		
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	-0.671	(-1.46)	-0.548	(-0.99)	1.501	(1.17)	0.781	(0.65)		
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-0.896	(-0.52)	-0.885	(-0.50)	-2.800	(-3.25)	-2.239	(-2.93)		

Panel A reports intraday expected variance risk premium components from non-overlapping 22-day windows. PCA factors are re-estimated each window from the cross-section of 5-minute stock returns (top 550 by market cap). Projection coefficients (θ) are estimated daily (within window) and lagged by one trading day. Stock returns are winsorized at 4σ per stock-day before PCA; factors are recomputed from original (non-winsorized) returns. Panel B reports overnight (close-to-open) components using rolling 22-day PCA windows. $K^\star$: Bai and Ng (2002) IC3 optimal; $2K^\star$: double $K^\star$. Annualized ($\times 76 \times 252$ for intraday, $\times 252$ for overnight). Newey and West (1987) t-statistics (lag = 21) in parentheses. The sample period is 2016–2025.

Table 2 reveals some differences between the intraday and overnight periods regarding the sources of priced risk. For intraday periods, the estimate for $\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$ appears small and statistically insignificant while the opposite is true for the overnight periods. This indicates that option markets demand premium for carrying overnight market variance risks that cannot be hedged dynamically by a portfolio of stocks. This result can be due to the fact that the weights in the hedge portfolio of stocks are kept fixed for a longer period when hedging overnight market variance risk than when hedging intraday market variance risk. This difference in frequency of rebalancing the hedge portfolio is due to the lack of high-frequency data during periods when markets are closed. An alternative explanation is that the population of investors actively setting prices in the options market overnight is dominated by intermediaries facing hedging frictions, and these intermediaries demand

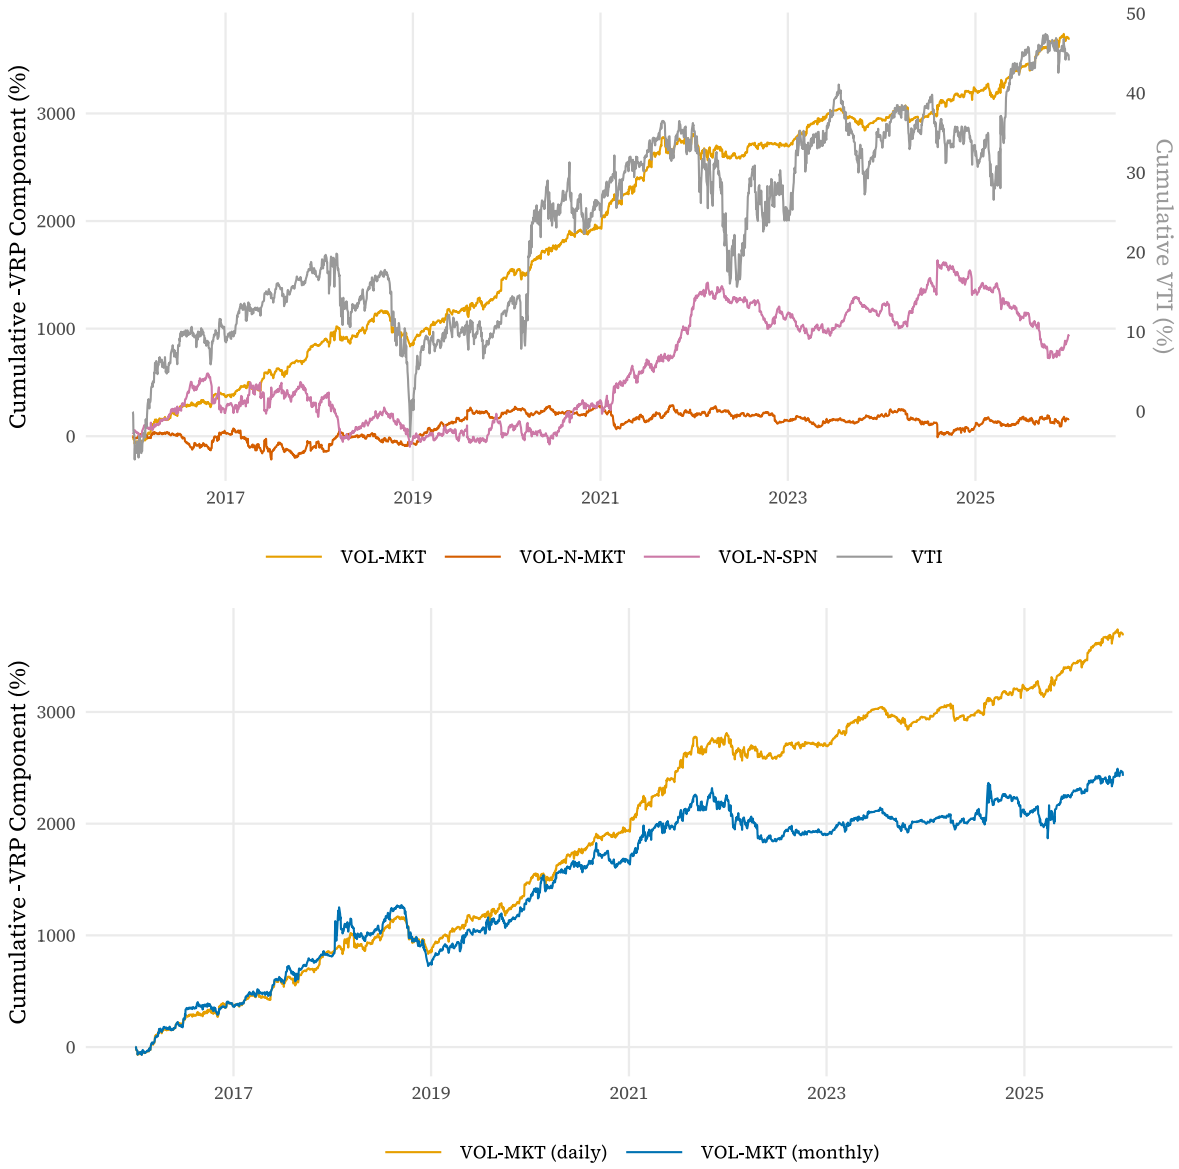
compensation for variance risks that they cannot hedge by trading stocks — which is precisely the risk captured by $\overline{R}_{s,t}^{\perp cs}$. We leave further investigation of this to future work.

We turn next to the expected market variance risks that are spanned by dynamic portfolios of stocks. Table 2 shows that effectively all priced risk of $\overline{R}_{s,t}^{cs}$ is attributed to $\overline{R}_{s,t}^{m}$, i.e. to the component that is dynamically spanned by trading only the market index, as $\mathbb{E}(\overline{R}_{s,t}^{\perp m})$ is statistically insignificant. Put differently, our evidence shows that all expected market variance risk that is priced in the cross-section of stock returns can be explained by the dependence of the latter on the market factor, i.e., by the so-called dynamic leverage effect. Once this dependence has been accounted for, there appears to be no role for expected market variance risk as a separate risk factor in explaining the cross-section of stock returns. This result holds true both for intraday and overnight market variance risk. It implies that the negative spread on volatility-beta-sorted portfolios documented by Ang et al. (2006) is not driven by a separately priced volatility factor. We will analyze the reasons for the differences in these findings in Section 6.

5.3 Cumulative Returns of the Components

In the top panel of Figure 2, we plot the cumulative returns of the options portfolio replicating expected market variance risk and its components. Consistent with the results reported in Table 2, the hedged option portfolio using the cross-section of stocks has a low average return. The cumulative return of $\overline{R}_{s,t}^{\perp cs}$ is negligible over the sample period and subperiods. The same holds true for the market-neutral portfolio of stocks that mimics expected market variance risk, i.e., for the cumulative $\overline{R}_{s,t}^{\perp \mathrm{m}}$. On the other hand, the dynamic portfolio of the market index that mimics the market variance risk yields consistently high returns over the entire sample. Comparison of the performance of this portfolio and the market index portfolio itself reveals nontrivial differences in their dynamics. This highlights the dynamic nature of the relation between market risk and expected market variance risk.

Figure 2: **Cumulative Intraday EVP Components**



The figure plots cumulative daily returns (in percent) of the negative intraday EVP and its components, 2016–2025. VOL-MKT ($\mathbb{E}(\overline{R}^{m}_{s,t})$) is the component of the PCA-spanned EVP explained by the market factor; VOL-N-MKT ($\mathbb{E}(\overline{R}^{\perp m}_{s,t})$) is the orthogonal, non-market component; VOL-N-SPN ($\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$) is the residual not spanned by the first $K^{\star}$ PCA factors. The bottom panel compares daily ($W = 1$) and monthly ($W = 22$) re-estimation of the spanning projection θ for VOL-MKT. All series are sign-flipped so that positive values correspond to a positive premium from selling variance exposure; cumulative VTI return (right axis) is shown in the top panel for reference.

6 Reconciling with Earlier Evidence for Pricing of Expected Market Variance Risk

Our finding that $\mathbb{E}(R_{s,t}^{\perp m})$ is statistically indistinguishable from zero — that expected market variance is not a separately priced risk factor for the cross-section of stocks — appears to contradict an influential body of empirical work, beginning with Ang et al. (2006) and continuing through Detzel et al. (2023), that argues for volatility risk being priced even after controlling for market-return exposure. To reconcile these findings, we begin by reproducing the Ang et al. (2006) portfolio-sorting exercise in our sample. Each month, we estimate stocks' volatility betas from daily regressions of returns on the market return and daily changes in the VIX, sort stocks into quintile portfolios on these betas, and compute the high-minus-low spread. We modify their procedure only by using intraday rather than daily data to improve the precision of the beta estimates and by using our model-free variance proxy in place of the VIX. The results, reported in Table 3, recover their finding: the high-minus-low quintile spread is negative and statistically significant (-0.72% per month, $t = -2.29$), and remains significant after adjusting for the CAPM and Fama-French (static) three-factor models. On our sample, then, the methodology of Ang et al. (2006) yields the conclusion that volatility risk is separately priced — directly contradicting the conclusion of our decomposition analysis in Section 4.

Table 3: Quintile Portfolios Sorted on Intraday Volatility Betas

Portfolio	Mean	Std	α_{CAPM}	β_{CAPM}	α_{FF3}	$\alpha^{\text{PCA}}_{K=3}$	$\alpha^{\text{PCA}}_{K=5}$	$\alpha^{\text{PCA}}_{K^*}$	$\bar{\beta}_V$	$\bar{\beta}^{\text{post}}_V$
1	1.52	4.96	0.38	1.00	0.34	-0.07	-0.04	-0.25	-0.0055	-0.0025
	(3.80)		(2.04)	(19.45)	(2.49)	(-0.39)	(-0.20)	(-1.52)		
2	1.01	4.65	-0.09	0.96	0.04	-0.18	-0.24	-0.43	-0.0024	-0.0017
	(2.95)		(-0.65)	(29.32)	(0.27)	(-1.38)	(-2.24)	(-3.25)		
3	0.88	4.44	-0.18	0.93	-0.00	-0.13	-0.24	-0.34	-0.0007	-0.0013
	(2.51)		(-0.99)	(19.63)	(-0.03)	(-0.81)	(-1.52)	(-2.29)		
4	1.11	4.81	-0.01	0.99	0.12	-0.09	-0.21	-0.27	0.0010	-0.0010
	(2.87)		(-0.09)	(24.87)	(0.88)	(-0.69)	(-1.37)	(-1.75)		
5	0.81	5.22	-0.33	1.00	-0.27	-0.44	-0.36	-0.34	0.0043	-0.0002
	(2.20)		(-1.61)	(15.18)	(-1.32)	(-2.60)	(-2.33)	(-2.44)		
5-1	-0.72	4.14	-0.71	-0.00	-0.61	-0.37	-0.32	-0.08		0.0023
	(-2.29)		(-2.12)	(-0.02)	(-1.71)	(-1.88)	(-1.44)	(-0.45)		

Quintile portfolios sorted on the pre-ranking intraday volatility beta $\beta^{\Delta \log V}$, estimated each month from OLS regressions of 5-minute log returns on the VTI return and $\Delta \log V$, where V is spot variance. Stocks are assigned to bins in month t using the beta from month $t-1$ and held for one month; portfolios are value-weighted. The "5-1" row is the high-minus-low spread. "Mean" and "Std" are the time-series mean and standard deviation of monthly portfolio excess returns (in percent); α_{CAPM} and β_{CAPM} are the CAPM intercept and market-beta; α_{FF3} is the Fama–French 3-factor intercept. $\alpha^{\text{PCA}}_{K=k}$ is the dynamically hedged alpha, $\alpha_p(t) = R_p(t) - \beta_p(t-1)'F(t)$, with $\beta_p(t-1)$ from prior-day OLS of 5-minute portfolio returns on 5-minute factor returns (the leading k WPCA factors from 22-day non-overlapping windows of 5-minute stock returns), summed within months. $K^\star$ is the per-window Bai and Ng (2002) IC3 optimal. $\bar{\beta}_V$ ($\bar{\beta}^{\text{post}}_V$) is the within-bin value-weighted average of the pre-formation (post-formation) $\beta^{\Delta \log V}$, time-series averaged. Newey and West (1987) t-statistics in parentheses (3 lags). The sample period is 2016–2025.

The central reconciliation result of this section is that this spread does not survive the relaxation of a single methodological assumption: the constancy of volatility betas over the monthly estimation window. Liao and Todorov (2026) reject this constancy formally, showing that the projection of variance shocks on systematic factors driving stock returns varies meaningfully at horizons shorter than a month. When we repeat the portfolio sort with weekly and daily rebalancing of the volatility-beta estimation — frequencies at which the time-variation in betas can be better accommodated — the high-minus-low spread vanishes. Table 4 shows the result: the weekly-rebalanced spread is -0.31% ($t = -0.76$) and the daily-rebalanced spread is +0.19% ($t = 0.44$), neither economically nor statistically distinguishable from zero. Related to this, the α's with respect to static factor models, which were significantly negative in Table 3, become insignificant when using dynamic factor models. The economic interpretation is straightforward: imposing a single beta on a relationship that varies within a period generates an apparent cross-sectional spread in volatility expo-

sure that does not reflect a genuinely priced source of risk. Once the dynamic structure is accommodated, the spread disappears.

The bottom panel of Figure 2 illustrates this misattribution directly. The figure plots the cumulative return of the dynamic market-index portfolio that mimics expected variance risk, $R^m_{s,t}$, under two different rebalancing frequencies for the projection coefficient ϕ^*_u: daily and monthly. The two cumulative-return series share a common upward trend — both portfolios earn the significant market-spanned variance risk premium documented in Table 2 — but they diverge in ways that are themselves informative. The monthly-rebalanced portfolio responds more strongly to large market moves, accumulating disproportionate gains and losses around episodes of elevated volatility. This is exactly the signature of a static specification picking up dynamic exposure: when the projection coefficient is held fixed across a month, the portfolio's response to any given market move reflects the average exposure over that month, missing the time-variation that could have moderated the response within higher-volatility subperiods. The same misattribution that distorts the volatility-mimicking portfolio under monthly rebalancing is at work in the Ang et al. (2006) cross-sectional regressions: in both cases, a static specification absorbs dynamic market exposure into what appears to be a separate volatility-related quantity.

In the last column of Table 3, we report the post-formation-period volatility betas. We note that the ranking of these betas across the sorted portfolios is preserved and is the same as the one of the formation-period volatility betas. This suggests that the long-short volatility portfolio is capturing indeed volatility risk. Naturally, since the estimation of dynamic volatility betas is rather difficult, the spread in post-formation-period volatility betas shrinks somewhat. Nevertheless, it remains highly statistically significant. We note also that the shrinkage of the spread in post-formation-period volatility betas is much smaller than the one reported in Ang et al. (2006). This illustrates the advantages of using intraday data for the estimation of stocks' volatility sensitivities.

Table 4: Quintile Portfolios Sorted on Intraday Volatility Betas with Shorter Rebalancing Window

Portfolio	Mean	Std	α_{CAPM}	β_{CAPM}	α_{FF3}	$\bar{\beta}_V$	$\bar{\beta}_V^{\text{post}}$
Panel A: Weekly rebalancing							
1	1.17	5.37	0.04	1.05	-0.02	-0.0099	-0.0025
	(2.74)		(0.19)	(21.13)	(-0.08)		
2	1.16	4.56	0.15	0.93	0.10	-0.0041	-0.0021
	(3.20)		(0.98)	(23.41)	(0.65)		
3	1.29	4.47	0.30	0.91	0.22	-0.0009	-0.0019
	(3.74)		(1.83)	(20.69)	(1.39)		
4	1.01	4.59	0.01	0.93	-0.02	0.0023	-0.0015
	(2.93)		(0.04)	(28.63)	(-0.14)		
5	0.87	5.56	-0.29	1.07	-0.22	0.0084	-0.0007
	(1.94)		(-1.13)	(15.72)	(-0.89)		
5-1	-0.31	4.37	-0.33	0.02	-0.21		0.0018
	(-0.76)		(-0.86)	(0.21)	(-0.52)		
Panel B: Daily rebalancing							
1	1.08	5.56	-0.10	1.09	-0.09	-0.0213	-0.0028
	(2.42)		(-0.38)	(14.14)	(-0.32)		
2	1.36	4.83	0.31	0.97	0.22	-0.0082	-0.0025
	(3.71)		(1.66)	(25.11)	(1.28)		
3	1.07	4.38	0.08	0.92	0.04	-0.0011	-0.0023
	(3.27)		(0.64)	(45.35)	(0.31)		
4	0.93	4.41	-0.04	0.90	-0.07	0.0060	-0.0020
	(2.93)		(-0.28)	(25.81)	(-0.47)		
5	1.27	5.18	0.22	0.97	0.26	0.0196	-0.0015
	(2.96)		(0.81)	(13.14)	(1.07)		
5-1	0.19	4.77	0.32	-0.12	0.34		0.0013
	(0.44)		(0.63)	(-0.86)	(0.72)		

Quintile portfolios sorted on the intraday volatility beta $\beta^{\Delta \log V}$, estimated from OLS regressions of 5-minute log returns on the VTI return and $\Delta \log V$, where V is spot variance. Panel A (weekly rebalancing): $\beta^{\Delta \log V}$ is estimated each week from that week's 5-minute returns. Portfolios are formed each Monday on the prior week's β and held one week. Panel B (daily rebalancing): same with daily estimation, prior-day β, and one-day holding. Portfolios are value-weighted. "5-1" is the high-minus-low spread. Daily portfolio excess returns are compounded within calendar month, and the spread is built from the monthly $P5$ and $P1$ series. "Mean" and "Std" are the time-series mean and standard deviation of monthly portfolio excess returns (in percent); α_{CAPM} and β_{CAPM} are the CAPM intercept and market-beta; α_{FF3} is the Fama–French 3-factor intercept. $\bar{\beta}_V$ ($\bar{\beta}_V^{\text{post}}$) is the value-weighted average of the pre-formation (post-formation) $\beta^{\Delta \log V}$, time-series averaged. Newey and West (1987) t-statistics in parentheses (3 lags) on the monthly series. The sample period is 2016–2025.

Another major difference between our approach and Ang et al. (2006) is that we use a

traded options portfolio to gauge the pricing of expected market variance risk. This avoids the need to estimate time-varying assets' sensitivities to volatility risk, which as argued above, is rather challenging. In addition, this direct approach allows us to identify the component of expected market variance risk that does not generate systematic risk in the asset prices (this is captured by $\overline{R}_{s,t}^{\perp cs}$) and explore its pricing. In Section 4.3, we showed that $\overline{R}_{s,t}^{\perp cs}$ can naturally appear in equilibrium models. Empirically, $\overline{R}_{s,t}^{\perp cs}$ is nontrivial and in the case of overnight returns, it requires compensation for by investors trading options.

How does presence of $\overline{R}_{s,t}^{\perp cs}$ affect inference for expected market variance risk premium? If there is a component in an observed risk factor that is unrelated to risk in the stocks, this leads to a compression of the estimated stocks' betas towards this factor. This is an attenuation bias akin to the one happening in the classical errors-in-variables problem. A compressed cross-sectional spread in betas towards the factor of interest (in our case expected market volatility), in turn, leads to inflated estimates of the risk premia for it when using standard linear asset pricing methods, see e.g., Breeden et al. (1989).

Ang et al. (2006) also construct a volatility-mimicking portfolio by projecting VIX changes on volatility-sorted portfolios of stocks and use this portfolio in Fama-MacBeth regressions. They report high in-sample correlation between changes in VIX and their volatility-mimicking portfolio, which will suggest negligible $\overline{R}_{s,t}^{\perp cs}$ in our notation. This high correlation estimate, however, is a substantially upward-biased estimate of the true correlation as it is based on a daily regression over a month and with 5 explanatory variables, see e.g., Cramer (1987). For example, even if there was no association between dependent and explanatory variables in a standard linear regression, because of the small sample for the volatility projection regression, we can have a correlation of around 50% between dependent variable (changes in VIX) and fitted value (volatility-mimicking portfolio returns).

A final major difference between our approach and that of Ang et al. (2006) is that our approach does not require a full specification of all systematic risk factors affecting stocks. Our only requirement is that, in the spirit of (a dynamic version of) the arbitrage pricing theory (APT), systematic risk factors are factors that drive cross-sectional variation in stock returns. We identify the systematic risk factors (up to a rotation) by local PCA analysis. A similar approach in the static case is proposed by Giglio and Xiu (2021). Giglio and Xiu (2021) show that when adopting such an approach, which is robust to omitted relevant factors, many proposed observable factors do not appear to have pricing ability or

the evidence for their pricing ability gets diminished nontrivially.

7 Robustness Checks

In this section we perform several robustness checks. First, we check the sensitivity of our results to using coarser frequency of rebalancing the volatility-mimicking portfolios based on the market portfolio and market-neutral position in the cross-section of stocks. Second, we check sensitivity to intraday sampling frequency for computing systematic factor portfolios and the decomposition of the expected variance return. We further check sensitivity of the results to the number of estimated systematic risk factors and to a real-time determination of these factors. Finally, we perform the analysis using an alternative volatility proxy constructed from the options data.

7.1 Coarser Hedging Frequency

We start with investigating the effect from changing the frequency at which we are allowed to rebalance the positions in the market portfolio and the stocks in the variance-mimicking portfolios. In Panel A of Table 5, we report results when using monthly rather than daily rebalancing of the hedging portfolios consisting of stocks or the market index. This robustness exercise is of independent interest because it speaks directly to one of the central methodological distinctions between our approach and that of Ang et al. (2006) and subsequent work: whether the projection of variance shocks onto systematic risks is treated as constant over monthly windows or allowed to vary at higher frequency. Comparing Panel B with our main results in Table 2, we see that coarsening the hedging frequency materially changes the decomposition. The premium attributed to the market-spanned component $\mathbb{E}(R^m_{s,t})$ shrinks substantially — from -3.75 to -2.47 in the K^* specification — while the premia for the components that are not spanned by the market dynamically ($\mathbb{E}(R^{\perp m}_{s,t})$ and $\mathbb{E}(R^{\perp cs}_{s,t})$) become larger in absolute magnitude and, in several specifications, statistically significant. In other words, when we handicap our estimator to use the same monthly hedging frequency that underlies the Ang et al. (2006) approach, we begin to recover the kind of result they report — namely, a significant premium for components of variance risk apparently orthogonal to the market. This shift is intuitive: a static hedge applied to a dynamic factor cannot capture

time-variation in the projection coefficient, so risk that is dynamically spanned by the market gets misattributed to the residual. Together, these findings make the dynamic-vs-static distinction central to our reconciliation of the empirical evidence in Section 6, where we show that applying daily rather than monthly rebalancing in the Ang et al. (2006) portfolio-sorting exercise eliminates the volatility-beta spread that the prior literature has documented.

Table 5: EVP Estimates using Coarser Sampling and/or Hedging Frequency

	Number of factors K							
	3		5		10		$K^\star$	
Panel A: 5-Minute Returns, Monthly Hedging								
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-4.888	(-4.54)	-4.888	(-4.54)	-4.888	(-4.54)	-4.888	(-4.54)
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-1.924	(-2.10)	-2.414	(-2.66)	-3.313	(-3.38)	-3.319	(-3.39)
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-2.333	(-2.80)	-2.395	(-2.83)	-2.456	(-2.86)	-2.470	(-2.87)
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	0.409	(0.98)	-0.019	(-0.07)	-0.858	(-2.38)	-0.848	(-2.21)
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-2.964	(-3.04)	-2.474	(-2.66)	-1.575	(-1.62)	-1.570	(-1.64)
Panel B: 10-Minute Returns, Daily Hedging								
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-4.843	(-4.51)	-4.843	(-4.51)	-4.843	(-4.51)	-4.843	(-4.51)
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.803	(-4.45)	-3.987	(-4.40)	-4.802	(-4.54)	-5.298	(-3.85)
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-3.622	(-5.22)	-3.643	(-5.20)	-3.732	(-5.24)	-3.783	(-5.27)
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	-0.181	(-0.45)	-0.344	(-0.67)	-1.070	(-1.52)	-1.516	(-1.38)
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-1.040	(-1.01)	-0.856	(-0.80)	-0.041	(-0.04)	0.455	(0.31)

Robustness checks for the intraday EVP decomposition. Panel A: projection coefficients (θ) estimated once from all 22 days of each window and applied to all days (monthly hedging, in-sample), using 5-minute returns. Panel B: same as Table 2 Panel A but using 10-minute returns (aggregated from the 5-minute panel). $K^\star$: Bai and Ng (2002) IC3 optimal. Newey and West (1987) t-statistics (lag = 21) in parentheses. The sample period is 2016–2025.

7.2 Coarser Intraday Sampling Frequency

We continue next with checking the sensitivity of our findings to the intraday sampling frequency. In the analysis thus far, we sample the options and the cross-section of stocks every five minutes during trading hours. Here, we redo the analysis using a coarser sampling frequency of ten minutes. The coarser sampling frequency should provide robustness to

potential market microstructure issues but will naturally make the estimation results less precise.

The results for the coarser sampling frequency and daily hedging are reported in Panel B of Table 5. Comparing these results with the ones reported in Tables 2 and 7, we see that our main findings remain overall unchanged. Mainly, for the intraday periods, the estimate of $\mathbb{E}(\overline{R}_{s,t}^{\perp cs})$ is small and statistically insignificant, and the same holds true in general for the one of $\mathbb{E}(\overline{R}_{s,t}^{\perp m})$.

7.3 Number of Systematic Risk Factors

Our procedure does not require knowledge of all systematic risk factors, but it does require an estimate of the linear space they span, and hence an estimate of their number. Our main results use the Bai and Ng (2002) information criterion to select K^*, the estimated number of factors. To check sensitivity to this choice, we report in the last columns of Table 2 results for $2K^*$ factors — i.e., for substantially over-extracted factor sets. We perform this check only for intraday returns at the five-minute sampling frequency, since reliable estimation of a larger number of factors requires both a wide cross-section and a high sampling frequency. Comparing the K^* and $2K^*$ columns, the estimates of all components of the decomposition change only marginally. In particular, the headline finding — that $\mathbb{E}(R_{s,t}^{\perp m})$ is small and statistically insignificant intraday while the market-spanned component carries the entire priced expected variance risk — is preserved across both specifications. This insensitivity to over-extraction is informative: had the K^* PCA missed a relevant systematic factor that loads on market volatility, doubling the number of factors would have moved variance risk from $R_{s,t}^{\perp cs}$ into $R_{s,t}^{cs}$ and altered the decomposition. The stability of the estimates therefore suggests that the first K* factors capture the systematic risks in the cross-section that are relevant for spanning variance risk, and that the residual $R_{s,t}^{\perp cs}$ reflects variance risk genuinely unspanned by the cross-section rather than variance risk loading on omitted factors.

7.4 Real-Time Implementable Factor Estimation

A potential concern with the main analysis is that the systematic risk factors are estimated within the same window over which the returns of the decomposition components are computed. In principle, this in-sample estimation could mechanically affect the decomposition:

the PCA factors are constructed to explain cross-sectional variance in the estimation window, and one might worry that they capture window-specific features of the data that spuriously load on contemporaneous variance shocks. To address this, we re-estimate Table 2 using factor loadings determined from the prior estimation window, so that the factors used in the projection are based exclusively on data observed before the projection period. This procedure is implementable in real time and removes any look-ahead from the factor estimation.

The results are reported in Table 6. Comparing them with the in-window estimates in Table 2, we see no meaningful differences. The headline findings — a significant intraday market-spanned premium with insignificant residuals, a significant overnight residual $\mathbb{E}(R_{s,t}^{\perp cs})$ — are unchanged in both sign and magnitude. This rules out the concern that our main results are driven by in-window factor estimation: the systematic risks in the cross-section that are relevant for spanning variance shocks are stable enough across consecutive 22-day windows that using lagged factor loadings produces essentially the same decomposition. As a corollary, the decomposition we report is implementable by an investor in real time, which is the relevant benchmark for thinking about it as a tradeable strategy.

Table 6: EVP Estimates using Out-of-Sample Systematic Factors

	Number of factors K									
	3		5		10		$K^\star$		$2K^\star$	
Panel A: Intraday										
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-4.631	(-4.19)	-4.631	(-4.19)	-4.631	(-4.19)	-4.631	(-4.19)	-4.631	(-4.19)
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.873	(-4.78)	-3.931	(-4.69)	-4.100	(-4.74)	-3.588	(-4.26)	-4.486	(-4.75)
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-3.910	(-5.77)	-3.931	(-5.72)	-4.002	(-5.74)	-4.021	(-5.74)	-4.066	(-5.71)
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	0.037	(0.13)	-0.001	(-0.00)	-0.098	(-0.22)	0.433	(1.01)	-0.420	(-0.63)
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-0.757	(-0.73)	-0.700	(-0.65)	-0.531	(-0.50)	-1.042	(-1.01)	-0.145	(-0.12)
Panel B: Overnight										
$\mathbb{E}(\overline{R}^{ev}_{s,t})$	-5.597	(-4.90)	-5.597	(-4.90)	-5.597	(-4.90)	-5.597	(-4.90)		
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.754	(-5.26)	-3.854	(-5.86)	-3.642	(-4.15)	-4.060	(-5.93)		
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-4.017	(-4.18)	-4.082	(-4.18)	-4.232	(-4.23)	-4.095	(-4.16)		
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	0.263	(0.43)	0.228	(0.30)	0.590	(0.68)	0.035	(0.05)		
$\mathbb{E}(\overline{R}^{\perp cs}_{s,t})$	-1.842	(-1.58)	-1.742	(-1.52)	-1.954	(-1.48)	-1.537	(-1.29)		

Identical specification to Table 2 ("EVP Estimates") except the spanning projection uses a predetermined factor basis: the previous window's loadings β_{w-1} (Panel A, non-overlapping 22-day windows), and the prior disjoint 22-day block's loadings (Panel B, rolling overnight). The factors therefore use only data available before the projection window and are implementable in real time. θ is still estimated daily and lagged exactly as in the main table. The first estimation window is dropped (Panel A: window 1 plus the first lagged day of each window; Panel B: the first 22 days). PCA factors are re-estimated each window from the cross-section of 5-minute stock returns (top 550 by market cap). Stock returns are winsorized at 4σ per stock-day before PCA; factors are recomputed from original (non-winsorized) returns. $K^\star$: Bai and Ng (2002) IC3 optimal; $2K^\star$: double $K^\star$ (Panel A only). Annualized ($\times 76 \times 252$ for intraday, $\times 252$ for overnight). Newey and West (1987) t-statistics (lag = 21) in parentheses. The sample period is 2016–2025.

7.5 Alternative Variance Proxy

We end with considering an alternative option-based measure of variance. The main advantage of using our options portfolio based approach is that we can study what shocks to V_t (recall (11)) are priced, including ones that are not spanned by systematic risks in the cross-section of the stock prices. V_t is an instantaneous measure of total expected market variance, including one generated by jumps. We can alternatively use a model-free measure of variance that is only due to the diffusive component of the price, i.e., a measure of $\sigma_t^\top \sigma_t$. The cost of doing this is that the option-based estimate of $\sigma_t^\top \sigma_t$ is not a traded portfolio (just like the

VIX index) and this means that we will not be able to study the pricing of all shocks to $\sigma_t^\top \sigma_t$ but only the pricing of those that are spanned by systematic risks in the cross-section of stocks. Nevertheless, this will allow us to determine whether a variance-mimicking portfolio, constructed from stocks, is a separate risk factor once its dynamic exposure to the market portfolio has been properly accounted for.

Our estimate for $\sigma_t^\top \sigma_t$ is the one proposed by Todorov (2019). This is an option-based measure, which is also behind the SPOTVOL index quoted by the CBOE options exchange. The estimate is given by

$$V_{t,T}^c(u) = -\frac{2}{u^2} \log |\mathcal{L}_{t,T}(u)|, \tag{42}$$

where

$$\mathcal{L}_{t,T}(u) = 1 - (u^2 + iu) \int_0^{F_{t,T}} \frac{e^{iu \log(K/F_{t,T})}}{K^2} P_{t,T}(K) dK - (u^2 + iu) \int_{F_{t,T}}^{\infty} \frac{e^{iu \log(K/F_{t,T})}}{K^2} C_{t,T}(K) dK. \tag{43}$$

If u is chosen sufficiently high, then we have

$$V_{t,T}^c(u) \approx (T - t) \times \sigma_t^\top \sigma_t,$$

and we refer to Todorov (2019) for a precise statement of this result. Since $V_{t,T}^c(u)$ is a nonlinear transformation of an option portfolio, it is not the price of a traded asset (just like the VIX index is not). By analogy to our forward variance measure $V_{t,T,2T}$, we can construct the diffusive forward variance

$$V_{t,T,2T}^c(u) = V_{t,2T}^c(u) - V_{t,T}^c(u), \tag{44}$$

and we have

$$V_{t,T,2T}^c(u) \approx T \times \sigma_t^\top \sigma_t. \tag{45}$$

We set u in a data-driven way. More specifically, we set u to

$$u_t = \sqrt{\frac{2 \log(2)}{V_{t,2T}}}, \tag{46}$$

where $V_{t,2T}$ is defined in (15). The integrals on the right-hand sides of (15) and (43) are

computed via Riemann sum approximation. The feasible estimate of $V^c_{t,T,2T}(u_t)$ is henceforth referred to as SV and its end-of-day values in our sample are displayed in Figure 1.

In Table 7, we report pricing results for SV. Since SV is not a traded portfolio, we only report estimates of $\mathbb{E}(\overline{R}^{cs}_{s,t})$ and its components $\mathbb{E}(\overline{R}^{m}_{s,t})$ and $\mathbb{E}(\overline{R}^{\perp m}_{s,t})$. Comparing these estimates with their counterparts for our portfolio-based replication of V_t reported in Table 2, we note that they are quite similar.

Table 7: EVP Estimates using the SV Proxy

	Number of factors K									
	3		5		10		$K^\star$		$2K^\star$	
Panel A: Intraday										
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.825	(-4.54)	-4.247	(-4.67)	-4.980	(-4.72)	-4.413	(-4.53)	-5.193	(-4.03)
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-3.975	(-5.70)	-3.999	(-5.70)	-4.040	(-5.67)	-4.064	(-5.68)	-4.134	(-5.70)
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	0.151	(0.43)	-0.248	(-0.54)	-0.940	(-1.33)	-0.349	(-0.57)	-1.059	(-0.98)
Panel B: Overnight										
$\mathbb{E}(\overline{R}^{cs}_{s,t})$	-3.863	(-4.75)	-3.898	(-4.83)	-3.773	(-4.46)	-4.135	(-5.81)		
$\mathbb{E}(\overline{R}^{m}_{s,t})$	-3.411	(-5.68)	-3.527	(-5.79)	-3.670	(-5.95)	-3.523	(-5.79)		
$\mathbb{E}(\overline{R}^{\perp m}_{s,t})$	-0.451	(-0.99)	-0.372	(-0.79)	-0.103	(-0.15)	-0.612	(-1.13)		

Same specification as Table 2, using the stochastic volatility (SV) variance proxy. Panel A reports intraday components from non-overlapping 22-day windows, with projection coefficients (θ) estimated daily and lagged by one trading day. Panel B reports overnight (close-to-open) components using rolling 22-day PCA windows. Since SV is not a traded portfolio, only the spanned component $\mathbb{E}(\overline{R}^{cs}_{s,t})$ and its market/non-market split are reported. $K^\star$: Bai and Ng (2002) IC3 optimal; $2K^\star$: double $K^\star$. Newey and West (1987) t-statistics (lag = 21) in parentheses. The sample period is 2016–2025.

8 Conclusion

This paper provides a model-free analysis of the pricing of shocks to expected market volatility. Our methodological contribution is the construction of a tradable portfolio of short-dated index options that spans short-term forward variance and whose return isolates shocks to expected market variance from price risk. This construction sidesteps the parametric assumptions of earlier work, the price jump contamination of realized-variance-based measures, and the static-beta restrictions of existing cross-sectional approaches. Because the portfolio is traded, the price of expected variance risk is identified directly from its excess returns.

We find that expected market variance shocks earn a significant negative premium, consistent with investors disliking high-volatility states. Decomposing the premium into three orthogonal components, however, reveals that essentially all of the compensation paid in the stock market is for the component of variance risk that is dynamically spanned by the market portfolio itself — a reflection of the dynamic leverage effect linking market returns and market variance, rather than evidence of a separately priced source of risk. The component of expected variance risk that is market-neutral but spanned by other systematic risks in stocks is essentially unpriced: not merely small, but statistically indistinguishable from zero. The one exception is overnight: the component of expected-variance risk orthogonal to the entire cross-section of stocks is unpriced intraday but earns a significant premium overnight.

Our reading of the above empirical findings is not that volatility is irrelevant for asset prices, but rather that its role operates through the dynamics of market exposure and market risk premia rather than as an additional priced factor in the stochastic discount factor for stocks. Our findings do not mean that investors do not care about changes in the investment opportunity set. Indeed, there can be many other state variables, in addition to expected market variance, that change the latter and that can be of importance to investors. Examples include low frequency components of volatility (as opposed to total volatility), tail risks, cross-sectional dispersion, etc. We leave the analysis of their pricing for future work.

9 Appendix

9.1 Implementation Details for the Expected Variance Risk Decomposition

The main empirical results of the paper, reported in Table 2, are based on estimations using non-overlapping rolling L-day windows of high frequency returns. Specifically, we use five-minute return data within each estimation window, where each window spans L trading days. For each window, we estimate the projection coefficient θ and perform the variance decomposition. The estimation window is then rolled forward by L trading days, and the estimation and decomposition procedures are repeated. As a result, the estimation windows are non-overlapping.

Below, we detail each step of the estimation procedure:

Step 1: Jump-winsorization. Let $R_{i,t}$ denote the simple return for firm i at period t, we winsorize the jump by applying

$$Y_{i,t} = (R_{i,t} \vee -4\sigma_i) \wedge 4\sigma_i$$

where σ_i denotes the five-minute stock's volatility computed from daily realized volatility. For the w-th estimation window, which spans L days, let

$$Y^{(w)} = (Y_{i,t}) \in \mathbb{R}^{N_w \times T_w}, \quad R^{(w)} = (R_{i,t}) \in \mathbb{R}^{N_w \times T_w}$$

be the jump-winsorized and non-winsorized return matrices, respectively, for N_w firms observed over the L-day estimation window, where T_w is the total number of intraday time intervals within the window.

Step 2: Weighted PCA (WPCA). Apply PCA on $Y^{(w)}$ to estimate K common components $\widehat{Y}^{(w)}$ and the idiosyncratic components $\widehat{E}^{(w)} = (\widehat{e}_{i,t}^{(w)}) = Y^{(w)} - \widehat{Y}^{(w)}$. The latter are used to estimate the idiosyncratic variance:

$$\widehat{\sigma}_i^2 = \frac{1}{T_w} \sum_{t=1}^{T_w} \left[\widehat{e}_{i,t}^{(w)}\right]^2.$$

Then, returns are re-weighted using the inverse of the idiosyncratic variance and PCA is applied again on the re-weighted returns to estimate the stock loadings in the systematic risk factors. Specifically, let $\widehat{\xi}$ be the $N_w \times K$ matrix of eigenvectors, corresponding to the large K eigenvalues of

$$W^{-1/2} Y^{(w)} Y^{(w)\prime} W^{-1/2},$$

where W is a diagonal matrix whose diagonal elements are $\widehat{\sigma}_1^2, ..., \widehat{\sigma}_{N_w}^2$. Then, estimate the factors as:

$$\widehat{F} = \frac{1}{\sqrt{N_w}} R^{(w)\prime} \widehat{\xi} \in \mathbb{R}^{T_w \times K}.$$

The number of factors K takes values from a predetermined set $K \in \{1, 5, 10, K^*, 2K^*\}$ where K^* is the number of factors selected using the information criterion of Bai and Ng (2002). We use the non-winsorized return to construct factors $\widehat{F}$ so that they can be treated as portfolio returns.

Step 3: Projection. Let $\Delta V^{(w)} \in \mathbb{R}^{T_w}$ denote the vector of changes in variances in the w-th estimation window. Then, regress the demeaned $\Delta V^{(w)}$ on demeaned $\widehat{F}$ to obtain the projection coefficient $\theta^{(w)}$. Specifically, let $M_T = I_{T_w} - 1_{T_w} 1'_{T_w}/T_w$,

$$\widehat{\theta}^{(w)} = (\widehat{F}' M_T \widehat{F})^{-1} \widehat{F}' M_T \Delta V^{(w)}.$$

Step 4: Variance Decomposition. The PCA spanned component of variance is then estimated by

$$\widehat{\Delta V}^{(PCA)} = \widehat{F}\widehat{\theta}^{(w)}.$$

A further regression of $\widehat{\Delta V}^{(PCA)}$ on the market return r^{mkt} decomposes $\widehat{\Delta V}^{(PCA)}$ into a market spanned piece and market-orthogonal piece:

$$\widehat{\Delta V}^{(PCA)} = P_{\text{mkt}} \widehat{\Delta V}^{(PCA)} + \widehat{e}, \quad P_{\text{mkt}} = r^{\text{mkt}} (r^{\text{mkt}\prime} r^{\text{mkt}})^{-1} r^{\text{mkt}\prime}.$$

Step 5: Subinterval aggregation. To mitigate measurement-error bias in raw option portfolio returns, we aggregate returns over subintervals $[s_i, t_i]$ within the reporting period $[s, t]$: Let $(A)_i$ denote the i th element of a vector A. We then have the following decomposition: (in the notation below, $\sum_i$ means the summation of all subintervals within a trading day which begins at period s and ends at period t.)

$$\begin{aligned} \underbrace{\frac{\sum_i (\Delta V^{(w)})_i}{\sum_i V_i^{(w)}}}_{\overline{R}^{ev}_{s,t}} = \underbrace{\underbrace{\frac{\sum_i (P_{\text{mkt}} \widehat{\Delta V}^{(PCA)})_i}{\sum_i V_i^{(w)}}}_{\overline{R}^{m}_{s,t}} + \underbrace{\frac{\sum_i (\widehat{\Delta V}^{(PCA)})_i - (P_{\text{mkt}} \widehat{\Delta V}^{(PCA)})_i}{\sum_i V_i^{(w)}}}_{\overline{R}^{\perp m}_{s,t}}}_{\overline{R}^{cs}_{s,t}} \\ + \underbrace{\frac{\sum_i (\Delta V^{(w)})_i - (\widehat{\Delta V}^{(PCA)})_i}{\sum_i V_i^{(w)}}}_{\overline{R}^{\perp cs}_{s,t}}, \end{aligned} \tag{47}$$

where each piece is intuitively described as:

$\overline{R}^{ev}_{s,t}$ — Total variance change
$\overline{R}^{cs}_{s,t}$ — PCA-spanned variance change
$\overline{R}^{m}_{s,t}$ — Market-spanned within $\overline{R}^{cs}_{s,t}$
$\overline{R}^{\perp m}_{s,t}$ — Market-orthogonal within $\overline{R}^{cs}_{s,t}$
$\overline{R}^{\perp cs}_{s,t}$ — Non-PCA spanned variance change

9.2 Proof of Proposition 1

Suppose $0 \leq s \leq t \leq T$. If a portfolio of options is started at time s that invests $1/K^2$ in out-of-the-money options with strike K, then the value of the portfolio at time t is given by

$$\int_0^{F_{s,T}} \frac{P_{t,T}(K)}{K^2} dK + \int_{F_{s,T}}^{\infty} \frac{C_{t,T}(K)}{K^2} dK, \tag{48}$$

while the value of a portfolio started at time t that invests $1/K^2$ in out-of-the-money options with strike K is given by

$$\int_0^{F_{t,T}} \frac{P_{t,T}(K)}{K^2} dK + \int_{F_{t,T}}^{\infty} \frac{C_{t,T}(K)}{K^2} dK. \tag{49}$$

The difference between the above two values is given by

$$DV_{s,t,T} = \begin{cases} \int_{F_{s,T}}^{F_{t,T}} \frac{C_{t,T}(K) - P_{t,T}(K)}{K^2} dK, & \text{if } F_{s,T} \leq F_{t,T}, \\ \int_{F_{t,T}}^{F_{s,T}} \frac{P_{t,T}(K) - C_{t,T}(K)}{K^2} dK, & \text{if } F_{s,T} > F_{t,T}. \end{cases} \tag{50}$$

Using put-call parity, this difference equals

$$DV_{s,t,T} = \begin{cases} \frac{1}{R_{f,t,T}} \int_{F_{s,T}}^{F_{t,T}} \frac{F_{t,T} - K}{K^2} dK, & \text{if } F_{s,T} \leq F_{t,T}, \\ \frac{1}{R_{f,t,T}} \int_{F_{t,T}}^{F_{s,T}} \frac{K - F_{t,T}}{K^2} dK, & \text{if } F_{s,T} > F_{t,T}. \end{cases} \tag{51}$$

Direct integration yields the following about the difference

$$DV_{s,t,T} = \frac{1}{R_{f,t,T}} \left[\frac{F_{t,T}}{F_{s,T}} - 1 - \log\left(\frac{F_{t,T}}{F_{s,T}} \right) \right]. \tag{52}$$

The above analysis means that the difference between the value of a portfolio at time t,

that is started at time s and invests $R_{f,s,2T}/K^2$ in out-of-the-money options with strike K and maturity $2T$ and shorts $R_{f,s,T}/K^2$ out-of-the-money options with strike K and maturity T and the value of the same portfolio of out-of-the-money options but started at time t is given by

$$\frac{R_{f,s,2T}}{R_{f,t,2T}}\left[\frac{F_{t,2T}}{F_{s,2T}}-1-\log\left(\frac{F_{t,2T}}{F_{s,2T}}\right)\right]-\frac{R_{f,s,T}}{R_{f,t,T}}\left[\frac{F_{t,T}}{F_{s,T}}-1-\log\left(\frac{F_{t,T}}{F_{s,T}}\right)\right]. \tag{53}$$

We note

$$\frac{R_{f,s,2T}}{R_{f,t,2T}}=e^{\int_s^{2T}r_u du-\int_t^{2T}r_u du}=e^{\int_s^{T}r_u du-\int_t^{T}r_u du}=\frac{R_{f,s,T}}{R_{f,t,T}}, \tag{54}$$

and

$$\frac{F_{t,2T}}{F_{s,2T}}=\frac{\widetilde{X}_t}{\widetilde{X}_s}\frac{R_{f,t,2T}}{R_{f,s,2T}}=\frac{\widetilde{X}_t}{\widetilde{X}_s}\frac{R_{f,t,T}}{R_{f,s,T}}=\frac{F_{t,T}}{F_{s,T}}, \tag{55}$$

where recall that the value of the market portfolio at time t is denoted with $\widetilde{X}_t$. Therefore, the expression in (53) is equal to zero. This shows that the value of the portfolio at time t, that is started at time s with value $V_{s,T,2T}$, equals $R_{f,s,t}V_{t,T,2T}$.

9.3 Proof of Proposition 2

We start with an expansion result for the realized variance risk premium. Using the expression for $DV_{s,t,T}$ in (52), we have the result in (18) for $R^{\text{rv}}_{s,t,T}$. Next, under standard smoothness in expectation assumptions for V_t, we have

$$\begin{aligned}\mathbb{E}_s^{\mathbb{P}}\left(\frac{V_{t,T}-V_{s,T}}{V_{s,T}}\right)&=\frac{T-t}{T-s}\frac{\mathbb{E}_s^{\mathbb{P}}(V_t)-\mathbb{E}_s^{\mathbb{Q}}(V_t)}{V_s}-\frac{t-s}{T-s}+O_P(T)\\&=\frac{T-t}{T-s}\frac{(t-s)EVP_s}{V_s}-\frac{t-s}{T-s}+O_P(T).\end{aligned} \tag{56}$$

Applying Ito's formula, we also have

$$\begin{aligned}2\mathbb{E}_s^{\mathbb{P}}\left[\frac{F_{t,T}}{F_{s,T}}-1-\log\left(\frac{F_{t,T}}{F_{s,T}}\right)\right]-(t-s)V_s&=(t-s)RVP_{s,t}+O_P(T^2)\\&=(t-s)RVP_s+O_P(T^2).\end{aligned} \tag{57}$$

Combining the above two results, we get the expansion result in (23).

We continue next with the expansion result for the forward variance risk premium. We have

$$\mathbb{E}_t^{\mathbb{Q}}\left(\int_{dT}^{dT+T} V_u du\right) = \mathbb{E}_t^{\mathbb{Q}}\left(\int_{dT}^{dT+T}\left(V_t + \int_t^u \alpha_z^{\mathbb{Q}} dz\right) du\right) = TV_t + \mathbb{E}_t^{\mathbb{Q}}\left(\int_{dT}^{dT+T}\int_t^u \alpha_z^{\mathbb{Q}} dz du\right), \tag{58}$$

and we have a similar expression for $\mathbb{E}_s^{\mathbb{Q}}\left(\int_{dT}^{dT+T} V_u du\right)$. From here, we can write

$$\begin{aligned}\frac{1}{T}\mathbb{E}_s^{\mathbb{P}}\big(V_{t,T,2T} - V_{s,T,2T}\big) &= (\mathbb{E}_s^{\mathbb{P}} - \mathbb{E}_s^{\mathbb{Q}})(V_t) \\ &\quad + \frac{1}{T}\mathbb{E}_s^{\mathbb{P}}\left[\mathbb{E}_t^{\mathbb{Q}}\left(\int_T^{2T}\int_t^u \alpha_z^{\mathbb{Q}} dz du\right) - \mathbb{E}_s^{\mathbb{Q}}\left(\int_T^{2T}\int_t^u \alpha_z^{\mathbb{Q}} dz du\right)\right],\end{aligned} \tag{59}$$

and we note that $(\mathbb{E}_s^{\mathbb{P}} - \mathbb{E}_s^{\mathbb{Q}})(V_t) = (\mathbb{E}_s^{\mathbb{P}} - \mathbb{E}_s^{\mathbb{Q}})(V_t - V_s)$. Next, we note that

$$\mathbb{E}^{\mathbb{P}}|\mathbb{E}_t^{\mathbb{Q}}(\alpha_z^{\mathbb{Q}}) - \mathbb{E}_s^{\mathbb{Q}}(\alpha_z^{\mathbb{Q}})| \leq CT, \tag{60}$$

for some constant $C > 0$. From here, we have

$$\frac{1}{T}\mathbb{E}_s^{\mathbb{P}}\left[\mathbb{E}_t^{\mathbb{Q}}\left(\int_T^{2T}\int_t^u \alpha_z^{\mathbb{Q}} dz du\right) - \mathbb{E}_s^{\mathbb{Q}}\left(\int_T^{2T}\int_t^u \alpha_z^{\mathbb{Q}} dz du\right)\right] = O_P(T^2). \tag{61}$$

From here, we get

$$\frac{1}{T}\mathbb{E}_s^{\mathbb{P}}\left(V_{t,T,2T} - V_{s,T,2T}\right) = (\mathbb{E}_s^{\mathbb{P}} - \mathbb{E}_s^{\mathbb{Q}})\big(V_t - V_s\big) + O_P(T^2), \text{ as } T \downarrow 0,\ 0 \leq s < t < T, \tag{62}$$

and hence the second result of the proposition.

References

Almeida, C. and G. Freire (2026). Demand in the option market and the pricing kernel. *Journal of Finance, forthcoming*.

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.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *The Journal of Finance 61*(1), 259–299.

Bai, J. and S. Ng (2002). Determining the number of factors in approximate factor models. *Econometrica 70*, 191–221.

Bakshi, G. and N. Kapadia (2003). Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies 16*(2), 527–566.

Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics 55*(2), 205–238.

Bates, D. S. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of econometrics 94*(1-2), 181–238.

Bollerslev, T., S. Z. Li, and V. Todorov (2016). Roughing up beta: Continuous versus 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. *The Review of Financial Studies 22*(11), 4463–4492.

Bollerslev, T. and V. Todorov (2011). Tails, fears, and risk premia. *The Journal of Finance 66*(6), 2165–2211.

Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger (1989). Empirical tests of the consumption-oriented capm. *The Journal of Finance 44*(2), 231–262.

Britten-Jones, M. and A. Neuberger (2000). Option prices, implied price processes, and stochastic volatility. *The journal of Finance 55*(2), 839–866.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley (2018). An intertemporal capm with stochastic volatility. *Journal of Financial Economics 128*(2), 207–233.

Campbell, J. Y. and T. Vuolteenaho (2004). Bad beta, good beta. *American Economic Review 94*(5), 1249–1275.

Carr, P. and D. Madan (2001). Optimal positioning in derivative securities. *Quantitative Finance 1*, 19–37.

Carr, P. and L. Wu (2009). Variance risk premiums. *The Review of Financial Studies 22*(3), 1311–1341.

Chen, H., S. Joslin, and S. X. Ni (2019). Demand for crash insurance, intermediary constraints, and risk premia in financial markets. *The Review of Financial Studies 32*(1), 228–265.

Cochrane, J. H. (2005). *Asset Pricing* (revised ed.). Princeton University Press.

Coval, J. D. and T. Shumway (2001). Expected option returns. *The journal of Finance 56*(3), 983–1009.

Cramer, J. S. (1987). Mean and variance of r2 in small and moderate samples. *Journal of econometrics 35*(2-3), 253–266.

Detzel, A. L., J. Duarte, A. Kamara, S. Siegel, and C. Sun (2023). The cross-section of volatility and expected returns: Then and now. *Critical Finance Review 12*, 9–56.

Dew-Becker, I., S. Giglio, and B. Kelly (2021). Hedging macroeconomic and financial uncertainty and volatility. *Journal of Financial Economics 142*(1), 23–45.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics 123*(2), 225–250.

Dickerson, A., M. Fournier, K. Jacobs, and P. Orłowski (2025). Systematic variance risk everywhere in equity option markets. *Available at SSRN 5741582*.

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

Duarte, J., C. S. Jones, and J. L. Wang (2024). Very noisy option prices and inference regarding the volatility risk premium. *The Journal of Finance 79*(5), 3581–3621.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. *The Journal of Finance 58*(3), 1269–1300.

Eraker, B. and I. Shaliastovich (2008). An equilibrium guide to designing affine pricing models. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics 18*(4), 519–543.

Eraker, B. and A. Yang (2022). The price of higher order catastrophe insurance: The case of vix options. *The Journal of Finance 77*(6), 3289–3337.

Garleanu, N., L. H. Pedersen, and A. M. Poteshman (2008). Demand-based option pricing. *The Review of Financial Studies 22*(10), 4259–4299.

Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy 129*(7), 1947–1990.

Hendershott, T., D. Livdan, and D. Rösch (2020). Asset pricing: A tale of night and day. *Journal of Financial Economics 138*(3), 635–662.

Jiang, G. J. and R. C. Oomen (2008). Testing for jumps when asset prices are observed with noise–a "swap variance" approach. *Journal of Econometrics 144*(2), 352–370.

Liao, Y. and V. Todorov (2026). Observable versus latent factors. *Journal of the American Statistical Association, forthcoming*.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society 41*(5), 867–887.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Ober, A. (2024). Intermediary inventory risk and the pricing kernel. *Available at SSRN 5032337*.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of financial economics 63*(1), 3–50.

Singleton, K. J. (2006). *Empirical dynamic asset pricing: model specification and econometric assessment*. Princeton University Press.

Todorov, V. (2019). Nonparametric spot volatility from options. *The Annals of Applied Probability 29*(6), 3590–3636.