Volatility and the Pricing Kernel*

David Schreindorfer
W.P. Carey School of Business
Arizona State University

Tobias Sichert
Department of Finance
Stockholm School of Economics

October 23, 2023

Abstract

We use options and return data to show that negative stock market returns are significantly more painful to investors when they occur in periods of low volatility. In contrast, asset pricing theories based on habits and long-run risks imply that the pricing of stock market risk does not vary with volatility, or that it moves in the opposite direction. An explanation of our finding that is consistent with prior empirical evidence and the time-varying disaster risk model of Gabaix (2012) is that volatility evolves independently of the pricing kernel.

JEL Classification: G12, G13, G33

Keywords: Pricing kernel, volatility, state prices, equity index options, tail risk, risk-return trade-off, conditional density estimation, habits, long-run risks, rare disasters.

*Contact: david.schreindorfer@asu.edu and tobias.sichert@hhs.se. We thank our discussants Ian Dew-Becker, Jens Christensen, Kris Jacobs, Christian Skov Jensen, Mete Kilic, Dmitriy Muravyev, and Paola Pederzoli, as well as Caio Almeida, Tyler Beason, Oliver Boguth, Adrien d'Avernas, Christian Julliard, Lars-Alexander Kuehn, Jonathan Payne, Seth Pruitt, Robert Ready, Guillaume Roussellet, Riccardo Sabbatucci, Petra Sinagl, Ivan Shaliastovich, Per Strömberg, and seminar participants at Arizona State, Stockholm School of Economics, McGill, Princeton, Carnegie Mellon, University of Oregon, Goethe University, ITAM, University of Georgia, University of Virginia, University of Iowa, Michigan State, and the University of Arizona, the Virtual Derivatives Workshop, JEF seminar, BI-SHoF Conference 2022, CICF 2022, SoFiE Conference 2022, EFA 2022, MFA 2023, and the SGF Conference 2023 for helpful comments and suggestions. Tobias Sichert acknowledges support from the Jan Wallanders and Tom Hedelius Foundation and the Tore Browaldhs Foundation, grant no. Fh21-0026.

Stock market risks are central for investors and policy makers alike. While time-variation in the quantity of stock market risk (e.g., volatility) is widely studied and well understood, the literature has devoted considerably less attention to time-variation in the pricing of stock market risks. In this paper, we show that the quantity and pricing of stock market risks are tightly connected. Specifically, we use index options and return data to show that negative stock market returns are significantly more painful to investors when they occur in periods of low volatility.

Our finding is reflected in the pricing kernel's projection onto stock market returns, which reveals how investors' marginal utility varies with returns. We propose a maximum likelihood estimator of the projection that allows us to condition on ex ante stock market volatility. Figure I illustrates our parametric estimate for the 10th and 90th percentile of volatility. The steeper curve in periods of low volatility shows that negative returns tend to coincide with significantly higher levels of marginal utility when they occur in calm markets. For example, monthly returns of -10% coincide with an average pricing kernel value of 3.68 when they occur in periods of low volatility, but only a value of 1.32 when they occur in periods of high volatility.

The time-varying pricing of stock market risks is also evident in option returns. Specifically, we use predictive regressions to show that stock market volatility is a positively and significantly related to returns on out-of-the-money put options. For example, puts that are struck 10% below the money earn an average return of nearly -100% in periods of low volatility, but only -70% in periods of high volatility. Like the pricing kernel evidence, this non-parametric finding implies that investors are more fearful of stock market crashes during calm markets.

To understand the economic forces behind our finding, it is essential to distinguish between the pricing kernel, M_{t+1}, which is potentially a function of many different shocks, and its projection onto stock market returns, $E_t[M_{t+1}|R_{t+1}]$, which is only a function of contemporaneous returns and investors' ex ante information set. The shape of the projected pricing kernel reflects the conditional joint distribution of returns and the pricing kernel. To illustrate the role of volatility, we derive

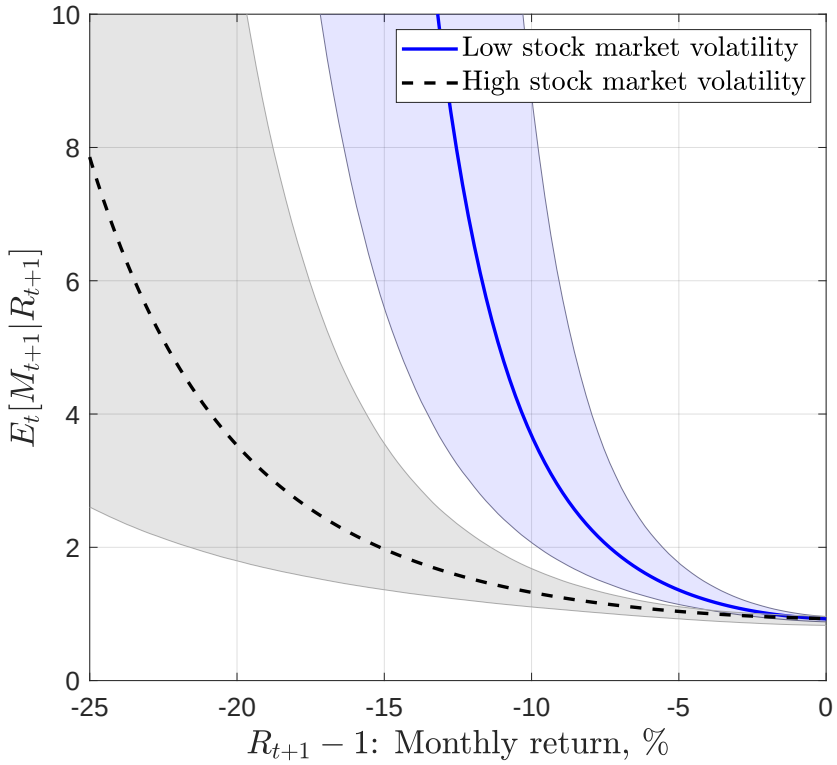

Figure I: Volatility and the projected pricing kernel. We plot the projected pricing kernel for the 10th and 90th percentile of conditional stock market volatility. The pricing kernel is measured at a monthly horizon, parameterized by equations (4) and (5) with a polynomial order of $N = 2$, and estimated over the 1990–2019 sample. Shaded areas represent pointwise 90% confidence bounds.

$E_t[M_{t+1}|R_{t+1}]$ under the assumption that R_{t+1} and M_{t+1} are conditionally jointly lognormal. This setting is stylized, but it encompasses the models of Campbell and Cochrane (1999) and Bansal and Yaron (2004), and we find that its main prediction carries over to models with non-normal shocks. Our "Proposition 1" shows that *a rise in stock market volatility makes the projected pricing kernel flatter, as in the data, if it is not accompanied by a rise in the conditional volatility of the pricing kernel.*

Building on this theoretical result, we show that our empirical finding can be used to distinguish between alternative economic mechanisms for time-varying stock market volatility. In contrast to Figure I, we find that the habit model of Campbell and Cochrane (1999) implies that negative returns are less painful in calm markets,

whereas the long-run risks model of Bansal and Yaron (2004) implies that investors are indifferent to the timing of negative returns. The same holds true for other models with recursive utility and persistent state variables, including Drechsler and Yaron (2011), Wachter (2013), and Constantinides and Ghosh (2017), whose implications for the projected pricing kernel closely mimic those of the original long-run risks model. We show that the counterfactual behavior of the projected pricing kernel in these models is an inherent feature of their economic mechanisms. In particular, both habit and long-run risks models rely on positive comovement between the conditional volatilities of returns and the pricing kernel to rationalize asset price dynamics, and therefore do not satisfy the condition of our Proposition 1. Without this positive comovement, neither model would be able to explain the countercyclical nature of volatility, the leverage effect of Black (1976), or the long-horizon predictability of excess returns, and the long-run risks model would lose a key source of the equity premium.

More encouragingly, we find that the model of Gabaix (2012) is consistent with our empirical evidence. In his model, stocks' exposure to potential consumption disasters varies over time and induces movements in stock market volatility that are independent of the pricing kernel. Different from models with habits or long-run risks, periods of high return volatility therefore do not coincide with elevated pricing kernel volatility. As predicted by our Proposition 1, and in line with the data, the projected pricing kernel in Gabaix's model therefore becomes flatter when return volatility rises. Another possible explanation of our finding lies in the "valuation risk" framework of Albuquerque et al. (2016). Their study augments the utility function of Epstein and Zin (1989) with shocks to agents' time preferences. These shocks result in large homoscedastic shocks to the pricing kernel and therefore add a time-invariant constant to the dynamics of its volatility. This constant implies that any changes in the pricing kernel's volatility are small on a relative basis. Hence, preference shocks provide an approximation to the case where pricing kernel volatility does not move with return volatility at all. Schorfheide et al. (2018) embed

this mechanism into a model with long-run risks. We find that the projected pricing kernel in their model behaves qualitatively as in the data, but the amount of time-variation is quantitatively much smaller than in the data. Valuation risk therefore appears to provide at best a partial explanation for our finding.

It is worth noting that Gabaix's model implies a form of independence that goes beyond the condition in Proposition 1. Specifically, his model predicts that return volatility is not only independent of pricing kernel's conditional volatility, but also independent of the pricing kernel realization itself. This is a subtle but important distinction. It is conceivable that shocks to return volatility affect the realized value of the pricing kernel without altering its conditional volatility. However, prior empirical evidence supports the stronger form of independence in Gabaix' model: Dew-Becker et al. (2017) show that shocks to expected (as opposed to realized) stock market volatility command no risk premium, while Berger et al. (2020) find that such shocks are not associated with macroeconomic contractions. It is therefore plausible that marginal utility (the realized value of M_{t+1}) is unaffected by news about stock market volatility. In our view, the collective evidence is most consistent with return volatility evolving independent of the pricing kernel's realization *and* its conditional distribution.

We build on a large literature that studies the pricing kernel's projection onto stock market returns based on index options, starting with Aït-Sahalia and Lo (2000), Jackwerth (2000), and Rosenberg and Engle (2002). The central finding of these studies is that the pricing kernel is a non-monotonic function of stock market returns – an observation dubbed the "pricing kernel puzzle" due to its inconsistency with standard models.[1] We contribute to this literature by documenting that the shape of the projection varies systematically with volatility.

In contemporaneous work, Kim (2022) examines time-variation in the projected

[1] In the online appendix, we show that our estimates are consistent with the projected pricing kernel's non-monotonicity. The main text focuses on the negative return region to draw attention to the novel fact we document – covariation with volatility – and away from the existing fact.

pricing kernel more broadly.[2] He considers several financial indicators as conditioning variables and finds results that align qualitatively with ours in the case of volatility. An important difference is that Kim follows the equity literature on conditional factor models, which commonly specifies betas as linear functions of conditioning variables (e.g., Nagel and Singleton 2011), and models the coefficients of his pricing kernel polynomial as linear functions of observables. We find that this approach substantially underestimates the amount of time-variation in $E[M|R]$ when volatility is used as a conditioning variable. Specifically, the non-linear coefficient specification we propose yields a significantly better fit to the data despite being more parsimonious. We further show that the much larger amount of time-variation implied by our specification is quantitatively consistent with non-parametric evidence from option returns.

Our study also makes a noteworthy methodological contribution. Specifically, we build on estimation approaches that take the risk-neutral distribution implied by option prices as given, parameterize the projected pricing kernel, and estimate parameters via a criterion function based on realized returns. Prior papers have implemented this idea in somewhat unconventional estimation frameworks: Bliss and Panigirtzoglou (2004) maximize the p-value of a Berkowitz test, whereas Linn et al. (2018) minimize a generalized method of moments criterion for moments of the inverse conditional CDF of returns. We follow the same general idea, but recognize that estimation can be performed via standard maximum likelihood. In particular, once a functional form has been specified for $E[M|R]$, it can be used to map the option-implied risk-neutral distribution to an analogous physical distribution, based upon which the likelihood function of returns can be computed. Doing so requires no additional parametric assumptions about the density of returns, which makes maximum likelihood a natural and elegant estimation approach in this context.

Finally, our focus on time-variation connects us to a large literature on the risk-

[2]Similarly, Driessen et al. (2020) examine whether the shape of the projected pricing kernel varies with the Chicago Fed National Activity Index, but find relatively small effects.

return trade-off in the time series of aggregate stock market returns. For example, the seminal study by Glosten et al. (1993) finds no significant relation between conditional volatility estimates from an asymmetric GARCH model and subsequent returns; Martin (2017) shows both theoretically and empirically that option-implied variance provides a (tight) lower bound for expected returns; Adrian et al. (2019) document a non-linear relation between the VIX index and expected returns. Hence, the results in this literature are fairly mixed. Our finding is consistent with an attenuated risk-return trade-off, as in Glosten et al. (1993), because it shows that times of high stock market volatility coincide with low risk prices (a flat projected pricing kernel). Yet our methodology reveals a more refined picture. It shows how individual state prices vary with volatility, whereas the aforementioned studies can only speak to the relation between volatility and expected market returns.

I. Estimating the Projected Pricing Kernel

This section explains our approach for estimating the pricing kernel as a function of stock market returns and conditional volatility, discusses data sources, and illustrates the robustness and statistical significance of our estimates. Throughout, the pricing kernel in period $(t+1)$ is denoted by M_{t+1}, the ex-dividend market return by R_{t+1}, and "t"- subscripts indicate moments and probability density functions that condition on investors' information set at time-t. We later drop 't"- subscripts for readability when not needed for clarity.

A. Estimation Approach

In the absence of arbitrage opportunities, the pricing kernel's projection onto stock market returns equals[3]

$$E_t[M_{t+1}|R_{t+1}] = \frac{1}{R_t^f}\frac{f_t^*(R_{t+1})}{f_t(R_{t+1})}, \tag{1}$$

where R_t^f is the risk-free rate and $f_t^*(R_{t+1})$ and $f_t(R_{t+1})$ denote the conditional risk-neutral and physical density of R_{t+1}, respectively. The projection measures the mean of M_{t+1} conditional on investors' information set at time-t and conditional on a (potential) return outcome at time-$(t+1)$. Apart from the market return, (1) therefore averages over all shocks that affect the pricing kernel at $(t+1)$. Importantly, (1) is generally a *nonlinear* conditional expectation function of R_{t+1} for any time-t information set, i.e., it is not a linear projection. Our estimation conditions on volatility as part of investors' time-t information set, as further detailed below.

To estimate $E_t[M_{t+1}|R_{t+1}]$, we extract f_t^* from option prices for each day of the sample based on the classic result of Breeden and Litzenberger (1978). This methodology is fairly standard and we refer interested readers to Appendix A for details. Next, we model the projection with a flexible parametric function of returns and the conditional return volatility, $M(R_{t+1}, \sigma_t; \theta)$, and combine it with (1) to

[3]The fact that the pricing kernel equals the ratio of risk-neutral to physical probabilities (scaled by R^f) is a well-known textbook result – see, e.g., Cochrane (2005), p. 51.

express the conditional physical density as

$$f_t(R_{t+1};\theta) = \frac{f_t^*(R_{t+1})}{R_t^f \times M(R_{t+1},\sigma_t;\theta)}. \tag{2}$$

Given a functional form for $M(R_{t+1},\sigma_t;\theta)$, the unknown parameter vector θ can be estimated by maximizing the log-likelihood of realized returns,

$$LL(\theta) = \sum_{t=1}^{T} \ln f_t(R_{t+1};\theta). \tag{3}$$

Our notation highlights f_t's dependence on the parameter vector θ, but it is important to emphasize that the density does not belong to a known parametric family of distributions. Rather, it results from applying a (parametric) change-of-measure to the (nonparametric) risk-neutral distribution f_t^*, whose shape is implied by the market prices of equity index options.

Our maximum likelihood estimator of $E[M|R]$ is statistically efficient and it incorporates conditioning information from the entire risk-neutral distribution. Both features represent important advantages over moment-based estimation approaches.

B. Parameterization

To ensure positivity, we model the projection as an exponential polynomial,[4]

$$M(R_{t+1},\sigma_t;\theta) = \exp\left\{\delta_t + \sum_{i=1}^{N} c_{it} \times (\ln R_{t+1})^i\right\}, \tag{4}$$

where the polynomial coefficients c_{it} vary with volatility according to

$$c_{it} = \frac{c_i}{\sigma_t^{b\times i}}, \tag{5}$$

and $\theta = (b, c_1, ..., c_N)$. Conditional volatility, σ_t, is estimated with the heterogeneous autoregressive (HAR) model of Corsi (2009) based on intradaily return data, as discussed in more detail in Appendix B.[5] The time-varying intercept, δ_t, is calculated

[4]Prior papers that have modelled the pricing kernel (or its projection onto returns) as a polynomial include Chapman (1997), Dittmar (2002), Rosenberg and Engle (2002), and Jones (2006).

[5]Figure IA.III of the online appendix shows that $E[M|R]$ displays even more time-variation when conditional volatility is emasured by the VIX index instead of a HAR model. This approach has the advantage of not relying on a parametric time series model for volatility, but the disadvantage that the VIX index embeds a risk premium and therefore systematically overestimates volatility.

for each day of the sample to satisfy the theoretical restriction that $f_t(R_{t+1};\theta)$ integrates to one, i.e., δ_t does not represent a free parameter.[6,7]

We experimented with different functional forms for the time-varying polynomial coefficients c_{it}, and found that (5) provides a very good fit (in terms of log-likelihood) despite its parsimony. Additionally, when we estimated a more flexible functional form for the relationship between c_{it}'s and σ_t, we found that its shape closely resembles the one in (5). This alternative specification for c_{it}'s is less parsimonious than (5) and we use it to illustrate the robustness of our results in Section I.F. Lastly, (5) nests two interesting special cases. For $b = 0$, the projected pricing kernel equals a time-invariant function of *returns*,

$$M(R_{t+1}, \sigma_t; \theta) = \exp\left\{\delta_t + \sum_{i=0}^{N} c_i \times (\ln R_{t+1})^i\right\}, \tag{6}$$

i.e., the graph of $E[M|R]$ does not vary with volatility, apart from a small vertical shift induced by δ_t. For $b = 1$, the projected pricing kernel equals a time-invariant function of *standardized returns* (up to a vertical shift due to δ_t),

$$M(R_{t+1}, \sigma_t; \theta) = \exp\left\{\delta_t + \sum_{i=0}^{N} c_i \times \left(\frac{\ln R_{t+1}}{\sigma_t}\right)^i\right\}. \tag{7}$$

In this case, the graph of $E[M|R]$ scales horizontally and proportionally with volatility. Intermediate values of b allow $E[M|R]$ to change with volatility to varying degrees. To formally test whether $E[M|R]$ varies with volatility, we evaluate the hypothesis $H_0 : b = 0$.

[6]The intercept equals $\delta_t = -\ln R_t^f + \ln\left(\int_0^\infty f^* \times \exp\left\{-\sum_{i=1}^{N} c_{it} \times (\ln R_{t+1})^i\right\} dR_{t+1}\right)$, i.e., its value is implied by R_t^f, f_t^*, and the polynomial coefficients $(b, c_1, ..., c_N)$. We find δ_t for each date by evaluating this integral numerically. By substituting the expression for δ_t into (4) and then (4) into (2), it can be verified that f_t integrates to one.

[7]Instead of computing δ_t based on the theoretical restriction $\int f = 1$, one could add a time-varying intercept c_{0t} to polynomial (4) and model c_{0t} as a function of volatility. Since this approach does not guarantee $\int f = 1$, however, it becomes necessary to add a penalty for violations of the restriction to the objective function. In turn, doing so requires the researcher to make a (necessarily subjective) choice on the relative importance of the restriction and the fit to realized returns. In the context of a moment-based estimation of the pricing kernel, Kim (2022) adds such a penalty term to his objective function, whereas Linn et al. (2018) simply ignore the theoretical restriction that probabilities need to add up to one.

C. *Parameter Identification*

The pricing kernel controls the extent to which conditional real world probabilities differ from their risk-neutral counterparts. Specifically, (2) shows that $f_t(R)$ takes on smaller values than $f_t^*(R)$ for return regions where $M(R_{t+1}, \sigma_t; \theta) > 1/R_t^f$ and higher values where $M(R_{t+1}, \sigma_t; \theta) < 1/R_t^f$. Individual elements of $\theta = \{c_1, ..., c_N, b\}$ are therefore identified if they alter the shape of $E[M|R]$ in such a way that it better explains the relative likelihood of different return realizations. Since f_t^* does not vary with θ, one can equivalently think of parameters as being identified by risk premia: An increase in the mean of f_t is equivalent to a higher equity premium, an increase in the variance of f_t is equivalent to a higher (less negative) variance premium, etc.

Most elements of θ alter the shape of f_t in multiple ways relative to that of f_t^*. Nevertheless, it is useful to discuss the main sources of parameter identification. c_1, the slope of $E[M|R]$, controls the relative probabilities of negative and positive returns. If the slope is negative, for example, the left tail of f_t^* gets downweighted in computing f_t, whereas the right tail gets upweighted. c_1 is therefore identified by the mean of f_t and the likelihood of negative returns. c_2, the curvature of $E[M|R]$, controls the relative probabilities of small and large absolute returns. If the curvature is positive, both extreme tails of f_t get downweighted relative to the tails of f_t^*, whereas the center of the distribution gets upweighted. Hence, c_2 is identified by the variance of f_t and the likelihood of extreme returns. c_2 is also negatively related to the mean of f_t because f_t^* is left-skewed, so that the equity premium further aids in its identification. Similarly, c_3, c_4, etc. are identified by higher order moments of f_t. The scaling parameter b controls how parameters of $E[M|R]$ vary with volatility, and therefore the amount of time-variation in the probabilities of different returns. For $b > 0$, an increase in volatility makes the slope of $E[M|R]$ less negative and its curvature less positive. b is therefore identified by the amount of time-variation in the moments of f_t, relative to time-variation in the corresponding f_t^* moments. We illustrate these channels quantitatively in Table IA.I of the online appendix by showing the sensitivity of moments of f_t to individual parameters.

D. Data

We use the S&P 500 index as a proxy for the aggregate stock market and focus on a return horizon of one month (30 calendar days). Return data comes from the Center for Research in Security Prices (CRSP). Price quotes on SPX options for the estimation of f_t^* come from the Chicago Board Options Exchange (CBOE). Because this data limits our sample to the 30-year period from 1990 to 2019, we sample daily to maximize the efficiency of our estimates, i.e., we work with a daily sample of $T = 7,556$ overlapping monthly returns. The estimation of conditional return volatilities (detailed in Appendix B) relies on intra-daily price quotes for S&P 500 futures, which were purchased from TickData. We use quotes for the large futures contract (ticker "SP") from 1990 to 2002, and for the E-Mini Futures contract (ticker "ES") from 2003 to 2019, i.e., we use data for the more actively traded futures contract in each part of the sample. Lastly, we use interest rates data from the Federal Reserve Bank of St. Louis' FRED database for robustness tests.

E. Estimation Results

Table I shows estimates for the parameterized pricing kernel in (4) and (5), and polynomial orders between $N = 1$ and $N = 5$. To account for autocorrelation that results from the use of overlapping return data, we determine the statistical significance of our estimates based on a block bootstrap with a block length of 21 trading days.[8]

The estimation results are easily summarized. The volatility-scaling parameter b is positive and significantly different from zero for all polynomial orders, and its significance grows in N. The observation that the shape of $E[M|R]$ varies with volatility is therefore not sensitive to the assumed polynomial order.[9] In fact, $E[M|R]$ is

[8] Volatility is also persistent, but this fact does not require a standard error adjustment because it does not induce autocorrelation into the observations that enter the objective function (3).

[9] One could additionally test whether the pricing of specific return states varies with volatility. We do so by bootstrapping the pointwise difference between $E[M|R]$ at the 10th and 90th percentiles

Table I: Estimation results

We estimate the projected pricing kernel in (4) and (5) for different polynomial orders N by maximizing the log likelihood of realized returns, (3). Statistical inference is based on a block bootstrap with a block length of 21 trading days. *, **, and *** denote significance at the 10%, 5% and 1% levels.

N	1	2	3	4	5
Log-likelihood	14,275	14,370	14,370	14,384	14,384
$\hat{c}_1$	-0.017^{***}	-0.067	-0.068	-0.095^{**}	-0.095^{**}
$\hat{c}_2$	–	0.100^{***}	0.103^{***}	0.020	0.020
$\hat{c}_3$	–	–	0.000	0.011^{*}	0.011^{*}
$\hat{c}_4$	–	–	–	0.003^{***}	0.003^{***}
$\hat{c}_5$	–	–	–	–	0.000
$\hat{b}$	1.600^{*}	0.976^{**}	0.973^{**}	1.098^{***}	1.097^{***}

well-described as scaling proportionally with volatility since the point estimate of b is close to one for all $N > 1$. As explained above, this implies that $E[M|R]$ is approximately a time-invariant function of standardized returns.

The log-likelihood increases substantially when the polynomial order is increased from $N = 1$ to $N = 2$, but only moderately thereafter. A likelihood ratio test rejects $N = 1$ in favor of $N = 2$ with a p-value of 0.15%, but fails to reject $N = 2$ in favor of any $N > 2$ at the 10% level (untabulated).[10] We show in Section I.G that the reason for the log-linear specification's poor fit lies in its inability to capture (i) the large amount of time variation in the projection and (ii) the large variance risk premium. Hence, the data clearly favors specifications for which the logarithm of $E[M|R]$ is convex. Since the parsimonious quadratic ($N = 2$) kernel is not rejected in favor of more flexible specifications, we use it as our benchmark case. All subsequent results are based on this estimate, unless otherwise mentioned.

Figure I in the introduction illustrates graphically how $E[M|R]$ varies with

of volatility. Figure IA.II of the online appendix shows that high volatility is associated with significantly lower risk prices (at the 1% significance level) for all monthly returns below -2%.

[10] There is no established method for dealing with overlapping data in likelihood ratio tests. We therefore rely on an ad-hoc sub-sampling approach. Specifically, we use observations $1, 22, 43, ...$, as the first subsample, observations $2, 23, 44, ...$, as the second subsample, and so on, up to observations $21, 42, 63, ...$, as the last subsample. We then estimate the two nested specifications of $E[M|R]$ in each subsample, compute their likelihood ratio, and average the individual likelihood-ratio statistics across the 21 subsamples. Finally, we compute critical values based on the statistic's asymptotic χ^2-distribution.

volatility by plotting it for the 10th and 90th percentile of σ_t (p_{10} and p_{90}). The figure shows that the pricing kernel is considerably steeper when volatility is low. For example, for a monthly return of -10%, the projected pricing kernel equals $M(R_{t+1} = -0.1, \sigma_t = p_{10}; \theta) = 3.68$ when volatility is low and $M(R_{t+1} = -0.1, \sigma_t = p_{90}; \theta) = 1.32$ when volatility is high.

Figure II shows the resulting conditional return densities for two dates. For comparability with Figure I, we choose days on which conditional volatility is close to its 10th and 90th percentile, respectively. Because our parameterization of $E[M|R]$ implies a smooth change-of-measure, f_t inherits many of f_t^*'s properties. It is unimodal, roughly bell-shaped, and its conditional volatility moves with that of f_t^*. Relative to f_t^*, however, f_t has more probability mass in the center and less mass in the left tail. As a result, the physical density is less left-skewed and leptokurtic than its risk-neutral counterpart, the equity premium is positive, and the variance premium is negative.

Across the 7,556 trading days in our sample, the conditional physical (risk-neutral) density has an average mean of 9.06% (0.98%) p.a., standard deviation of 13.83% (17.97%) p.a., skewness of -0.61 (-1.48), and kurtosis of 4.43 (10.47). The time series of these moments in Figure IA.IV of the online appendix shows that they are well-behaved. Our density estimates imply that the conditional equity premium $E_t[R_{t+1}] - E_t^*[R_{t+1}]$ has an average of 8.1% p.a., which closely matches the average excess return on the S&P 500 of 8.0% over the 1990–2019 period. Similarly, our density estimates imply that the conditional variance premium $var_t[R_{t+1}] - var_t^*[R_{t+1}]$ has an average of $-12.4\%^2$ per month, which closely matches the average $\sigma_t^2 - \left(\frac{VIX_t}{100}\right)^2$ of $-13.2\%^2$ per month over 1990–2019. The parametric pricing kernel therefore provides a good fit for stock market risk premia in our sample.

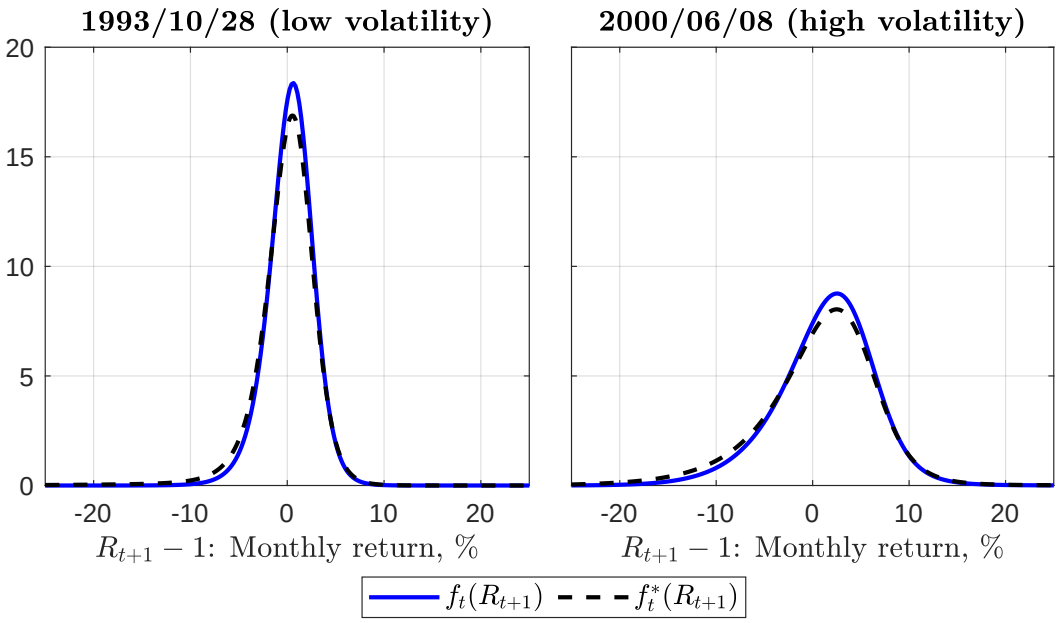

Figure II: Conditional density estimates for select days. We plot the estimated physical and risk-neutral return density for days on which conditional volatility is close to its 10th (left panel) or 90th (right panel) percentile. Estimates are based on the $E[M|R]$ specification in equations (4) and (5) and a polynomial order of $N = 2$.

F. Robustness

We perform four robustness tests. First, we model the projected pricing kernel's volatility-dependence with the alternative specification

$$c_{it} = \sum_{k=0}^{K} c_{ik} \times \sigma_t^k, \tag{8}$$

which assumes that coefficients of the $E[M|R]$-polynomial are themselves polynomials of volatility. The combination of (4) and (8) is equivalent to a bivariate polynomial in $\ln R_{t+1}$ and σ_t with a tensor product base. We find that, for $K = 2$ and higher orders, the estimated functional relationship between σ_t and c_{it}'s implied by (8) closely resembles the one in our benchmark specification (5). As a result, the relationship between σ_t and the shape of $M(R_{t+1}, \sigma_t; \theta)$ also closely resembles the one in our benchmark specification. We illustrate this fact for $K = 3$ below and for other polynomial orders in Section I.B of the online appendix. In contrast,

our estimate of the linear case ($K = 1$) implies a poor fit to the data, despite relying on an additional degree of freedom relative to our benchmark specification.[11] Specifically, we rely on Vuong's (1989) likelihood-ratio test for non-nested models to test the null hypothesis that our benchmark specification (5) and the $K = 1$ case of (8) are equally close to the true model against the alternative hypothesis that our benchmark specification is closer. The test rejects the null with a p-value of 0.41%. As we illustrate in the next section, this poor fit stems from the fact that the $K = 1$ case of (8) generates little time-variation in $E[M|R]$.

Second, we orthogonalize σ_t with respect to proxies of option market liquidity. This test addresses the potential concern that movements in $E[M|R]$ reflect frictions in the options market that are positively correlated with stock market volatility. We proxy option market liquidity by (1) daily option volume (across all maturities and strike prices), normalized by the average daily volume over the prior three month to remove the time trend in volume (following Chen et al. 2019), (2) daily open interest (across all maturities and strike prices), similarly normalized by the trailing three months average and (3) the bid-ask spread of the put option that is closest to being at-the-money and a maturity of 30 calendar days, normalized by the option's midquote. We then regress σ_t on the three liquidity variables, scale the resulting regression residuals so that they have the same mean and variance as σ_t, and re-estimate our benchmark specification of $E[M|R]$ based on this liquidity-adjusted volatility measure.

Third, instead of modelling $E[M|R]$ as a polynomial, we model $f_t(R_{t+1};\theta)$ as a parametric density and obtain $E[M|R]$ from (1) as the ratio of risk-neutral and physical densities, scaled by the risk-free rate. Specifically, we parameterize the density of standardized log returns, $g_t\left(\frac{\ln R_{t+1}}{\sigma_t};\theta\right)$, with a normal inverse Gaussian (NIG) distribution and compute the distribution of simple returns via a change of variables as $f_t(R_{t+1};\theta) = g_t\left(\frac{\ln R_{t+1}}{\sigma_t};\theta\right)/(\sigma_t \times R_{t+1})$. The NIG distribution is uni-

[11] The $K = 1$ case corresponds to the specification in Kim (2022), apart from the fact that he models the pricing kernel's intercept as parametric function of volatility, whereas we choose it such that the implied return densities integrate to one – see footnote 7 for details.

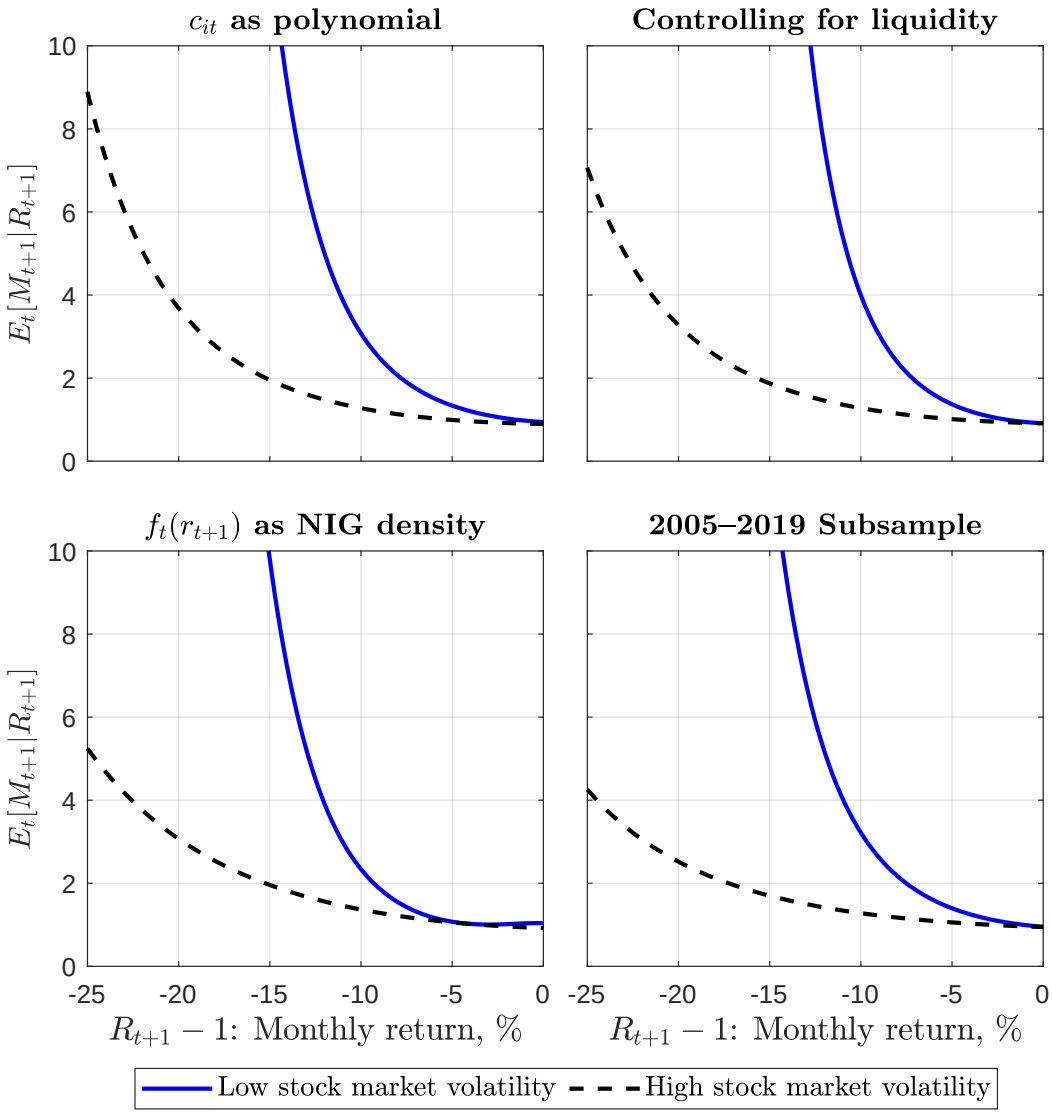

Figure III: Robustness. We plot the projected pricing kernel for the 10th and 90th percentile of conditional stock market volatility. Top-left: $E[M|R]$ is a $N = 2$ polynomial with coefficients that depend on volatility via (8) with $K = 3$. Top-right: $E[M|R]$ is equivalent to the benchmark specification, but σ_t is orthogonalized with respect to different liquidity measures. Bottom-left: We model the distribution of standardized log returns $\ln R_{t+1}/\sigma_t$ with a Normal Inverse Gaussian distribution, compute $f_t(R_{t+1})$ via a change-of-variables, and obtain $E[M|R]$ from (1). Bottom-right: $E[M|R]$ is equivalent to the benchmark specification, but estimated over the 2005–2019 subsample.

modal, bell-shaped, allows for nonzero skewness and excess kurtosis, and depends on four parameters, which we estimate via maximum likelihood. This method for estimating the conditional distribution resembles the popular approach of scaling historical return innovations from a GARCH model with a current estimate of conditional volatility – see, e.g., Rosenberg and Engle (2002), Barone-Adesi et al. (2008) and Christoffersen et al. (2013) – and shares its limitation that higher conditional moments (beyond volatility) are time-invariant by construction. In contrast, the parameterized pricing kernel in our benchmark specification allows all return moments to vary over time.

Fourth, we re-estimate the benchmark specification in the second half of the sample (2005–2019) to address concerns about a possible segmentation between index option and equity markets. In particular, Dew-Becker and Giglio (2022) argue that the two markets have historically been segmented, but also provide evidence suggesting that they have become well-integrated since about the mid 2000's. If time-variation in the estimated projected pricing kernel was a result of market segmentation, one would expect it to be substantially weaker in more recent data.

Figure III shows that, for each of the four alternative estimates, the projected pricing kernel's volatility-dependence looks similar to our benchmark estimates in Figure I. Parameter estimates for these specifications are reported in Section I–II of the online appendix. In the bivariate polynomial specification, a likelihood ratio test strongly rejects the hypothesis $H_0 : c_{ik} = 0 \ \forall i, k > 0$ (time-invariance) with a p-value of 0.2%. In the specification with liquidity controls, the estimated volatility-scaling parameter of $\hat{b} = 1.02$ is very close to the benchmark estimate of 0.976, and it remains statistically significant with a p-value of 1.3%. The parameterized density approach does not lend itself to a formal statistical test of whether or not the shape of $E[M|R]$ varies over time, but the amount of time-variation in $E[M|R]$ is quantitatively similar to that in Figure I. In the 2005–2019 estimation, the corresponding point estimate of b is once again similar to the benchmark at $\hat{b} = 1.01$. Our main result is therefore not sensitive to the way $E[M|R]$ is parameterized, the volatility-

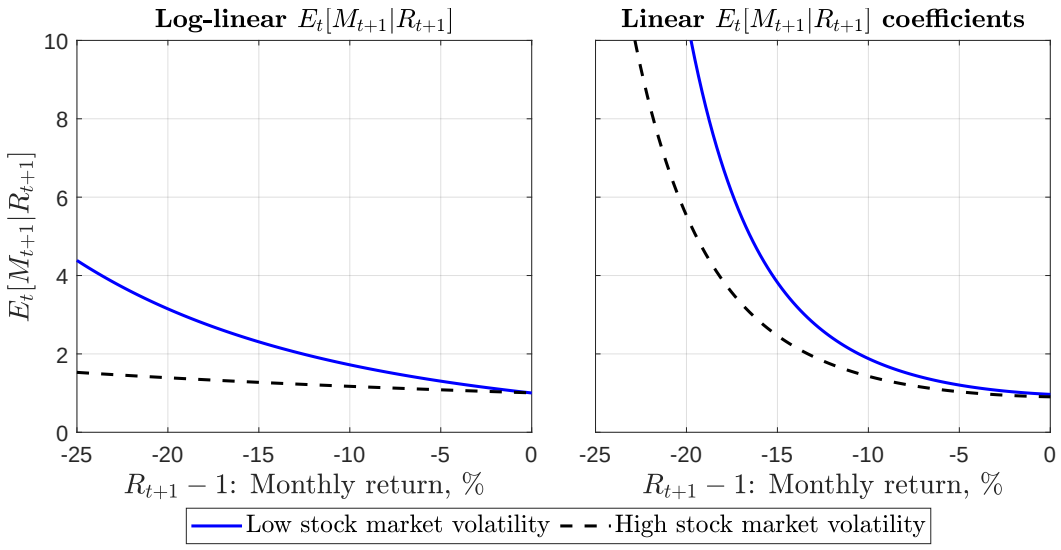

Figure IV: Alternative parameterizations of $E[M|R]$. We plot the equivalent to Figure 1 for a log-linear ($N = 1$) polynomial with slope coefficient $c_{1t} = c_1/\sigma_t^b$ (left panel) and for a log-quadratic polynomial with coefficients $c_{it} = c_{i0} + c_{i1} \times \sigma_t$ (right panel).

dependence of $E[M|R]$ does not reflect comovement between volatility and liquidity in the options market, and it can also not be explained by market segmentation.

G. Specification Issues

Our benchmark specification models both $E[M|R]$ and the volatility-dependence of its coefficients as non-linear functions. Is this necessary? We saw in Section I.E that a log-likelihood test strongly rejects the simpler log-linear polynomial for $E[M|R]$ in favor of our benchmark, whereas Section I.F showed the same result for a log-quadratic polynomial whose coefficients depend linearly of volatility. As we now show, an important reason for the statistical rejection of these simpler alternatives lies in their inability to capture the large amount of time-variation in $E[M|R]$.

Figure IV shows $E[M|R]$ for the 10th and 90th percentiles of volatility for the two alternative specifications. Clearly, the projected pricing kernel varies considerably less with volatility than for our benchmark parameterization in Figure 1. For example, for a return of -10%, the $N = 1$ specification in the left panel implies that average marginal utility decreases from 1.72 to 1.17 when moving from the 10th to

the 90th percentile of volatility and the log-quadratic case with linear coefficients in the right panel implies that it decreases from 1.88 to 1.43, compared to a decrease from 3.68 to 1.32 for our benchmark specification.

Three observations suggest that the large amount of time-variation in our benchmark specification is closer to the true data generating process. First, the benchmark provides a superior statistical fit to the data. Second, our first robustness test in Section I.F showed that an alternative specification which models $E[M|R]$'s coefficients as flexible higher-order polynomials of volatility generates nearly identical time-variation as the benchmark. Third, as we will show in the following section, the benchmark specification of $E[M|R]$ implies time-variation in expected option returns that aligns quantitatively with non-parametric evidence from realized option returns. In contrast, Figure IA.V in the online appendix shows that the simpler linear specifications imply expected option returns that do not align with the non-parametric evidence. They are therefore misspecified in the sense that they fail to capture most of $E[M|R]$'s time-variation.[12]

[12] Another clear indication of misspecification of the log-linear case lies in its inability to capture risk premia on higher return moments. For example, it implies that the average variance premium equals $-4.4\%^2$ per month with a 99% confidence interval of $[-7.1\%^2, -1.0\%^2]$. The $-13.2\%^2$ sample average of $\sigma_t^2 - \left(\frac{VIX_t}{100}\right)^2$ has a bootstrapped p-value of 0.00% under the sampling distribution of the $N = 1$ estimator. The reason for this mismatch can be gleaned from the left panel of Figure IV. Relative to the log-quadratic case, the log-linear pricing kernel is substantially flatter, especially in the far left tail of the return distribution. It is therefore also misspecified in the sense that it understates investors' aversion against tail events.

II. Volatility and Expected Put Option Returns

The time-varying pricing of stock market risks is also reflected in option returns. To see this connection, note that the physical return density $f_t(R_{t+1})$ can be interpreted as the expected payoff of an Arrow-Debreu (AD) security that pays \$1 if the realized stock market return equals R_{t+1}, and \$0 otherwise, whereas the discounted risk-neutral density $\frac{1}{R_t^f} f_t^*(R_{t+1})$ can be interpreted as the AD security's price. The projected pricing kernel $E_t[M_{t+1}|R_{t+1}] = \frac{1}{R_t^f} \frac{f_t^*(R_{t+1})}{f_t(R_{t+1})}$ equals the ratio of price-to-expected payoff, i.e., it measures the inverse of the AD security's expected return. A rise in stock market volatility decreases $E_t[M_{t+1}|R_{t+1}]$ and therefore increases the expected returns of AD securities.

To make the connection to option returns, note that a put option with moneyness K is a portfolio of AD securities for return states below K. Since the AD securities' expected returns rise when volatility increases, the put's expected return rises as well. To quantify this effect, we use the density estimates from our parametric estimation of $E_t[M_{t+1}|R_{t+1}]$ to compute expected put option returns,

$$E_t[R_{t+1}^{put(K)}] = \frac{\int_0^\infty f_t(R) \max\{0, K - R\} dR}{\frac{1}{R_t^f} \int_0^\infty f_t^*(R) \max\{0, K - R\} dR}, \tag{9}$$

for the 10th and 90th percentile of stock market volatility. Figure V shows that the resulting differences in expected returns are economically large. For example, protection against a -10% drop in the market is associated with an expected return of -96% when volatility is low (solid line), but only -70% when volatility is high (dashed line).

The positive association between volatility and expected option returns may be surprising to some readers, because it is well-known that options become more valuable when volatility rises. All else equal, higher option prices should result in lower expected returns. Yet our estimates of the projected pricing kernel imply that the opposite relation holds true in the data. This finding implies that expected option payoffs increase even more than option prices in times of high volatility, i.e.,

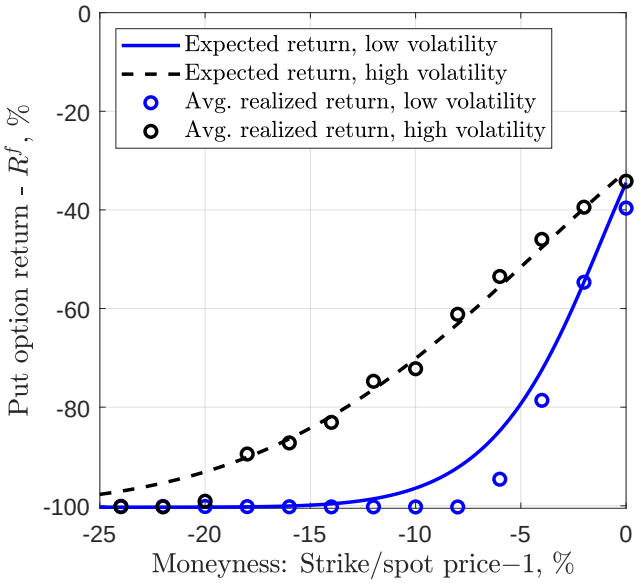

Figure V: Put option returns. The plot shows expected and average realized monthly S&P 500 put option returns. Expected returns are computed based on our parametric estimate of $E[M|R]$ and evaluated at the 10th and 90th percentile of σ_t, respectively. Average realized returns are computed separately for months in which the ex ante σ_t falls in its first or fourth quartile, respectively.

options are expensive in absolute (dollar) terms, but cheap in relative (return) terms in turbulent markets.

To address the concern that a potential misspecification of $E[M|R]$ leads to biased estimates of expected option returns, we compare the aforementioned estimates to average realized put returns. Specifically, we construct a panel of realized put returns with target moneyness levels of -24%, -22%,-20%,...,0%. For each month, we select a put option with 30 days to maturity that expires on the main expiration cycle (third Friday), whose moneyness is closest to (but is at most $\pm 1\%$ away from) the target, and compute its hold-until-maturity return. Lastly, we average returns separately for low and high volatility periods, which we define as months in which the ex ante σ_t falls in its first or fourth quartile, respectively. The circles in Figure V show that the magnitudes of these (nonparametric) average realized put returns closely mimic their (parametric) expected return counterparts.

To formally test whether expected option returns vary with volatility, we run

Table II: Predicting put returns with ex ante volatility

The table shows β regression coefficients from (10) for different moneyness levels. *, **, and *** denote significance at the 10%, 5% and 1% levels. Newey and West (1987) t-statistics with five lags are in parenthesis.

$K/S-1,\%$	-20	-18	-16	-14	-12	-10	-8	-6	-4	-2	0
$\hat{\beta}$	0.02	0.24	0.34^{**}	0.43^{**}	0.45^{**}	0.49^{**}	0.49^{**}	0.52^{**}	0.41^{*}	0.29	0.05
	(1.27)	(1.63)	(2.01)	(2.17)	(2.36)	(2.19)	(2.45)	(2.50)	(1.86)	(1.31)	(0.25)

the predictive regression

$$R^{Put}_{t+1,i} - R_{f,t} = \alpha + \beta \ln(\sigma_t) + \gamma K_i/S_t + \epsilon_{t+1,i}. \tag{10}$$

Since Figure V suggests that the interaction between moneyness and volatility is highly nonlinear, we run the regression separately for each of the aforementioned moneyness bins. We use a log-transformation of volatility because it is better-behaved, and control for moneyness to account for any remaining differences in expected returns due to differences in moneyness within bins.

Table II confirms that volatility predicts put returns with a positive sign in all moneyness bins, and does so significantly for moneyness levels between -16% and -4%.[13] Standard errors are adjusted for heteroscedasticity and autocorrelation based on Newey and West (1987) with five lags.

Overall, the evidence on time-variation in expected option returns confirms the finding that negative stock market returns are more painful to investors in times of low volatility. In addition to being non-parametric, option returns have the advantage of casting investors' willingness to pay for crash insurance in returns units, which are easier to interpret economically than marginal utility units.

[13] Hu and Jacobs (2020) and Aretz et al. (2022) find a similar relationship between firm-level volatility and returns of puts written on individual stocks. Their findings reflect cross-sectional variation in firm-level volatility, while our result reflects time-series variation in aggregate stock market volatility.

III. Economic Interpretation

Which economic forces cause time-variation in $E[M|R]$? What does the time-varying pricing of stock market risks reveal about the dynamics of volatility and its co-movement with M_{t+1}? What connections, if any, are there to other prominent asset market facts, such as the average equity premium, return predictability, or the countercyclical nature of volatility? This section attempts to answer these questions by examining the projected pricing kernel in a number of structural asset pricing models. Section III.A previews results for the models and illustrates our main theoretical result, whereas Sections III.B–E show how time variation in $E[M|R]$ relates to the models' economic mechanisms. Computational details are relegated to the online appendix.

A. Intuition and Overview of Results

Figure VI shows $E[M|R]$ in the models of Campbell and Cochrane (1999), Bansal and Yaron (2004), Gabaix (2012), and Schorfheide et al. (2018), while a number of additional models are analyzed in the online appendix and briefly discussed below. In contrast to the data, the habit model (top-left panel) implies that the projection becomes steeper when volatility rises, whereas the long-run risks model (top-right panel) implies that its shape does not vary with volatility. In line with the data, the time-varying disaster risk model (bottom-left panel) and valuation risk model (bottom-right panel) imply that the projection becomes flatter when volatility rises. To understand the source of this behavior and its link to the models' economic mechanisms, we first characterize the shape of $E_t[M_{t+1}|R_{t+1}]$ in a general lognormal setting, which nests three of the aforementioned models.

PROPOSITION 1. Assume that the log pricing kernel and log return follow

$$\begin{aligned}\ln M_{t+1} =& \mu_{m,t} - \sigma_{m,t}\varepsilon_{s,t+1}\\ \ln R_{t+1} =& \mu_{r,t} + \sigma_{s,t}\varepsilon_{s,t+1} + \sigma_{i,t}\varepsilon_{i,t+1},\end{aligned}$$

where the systematic shock $\varepsilon_{s,t}$ and the idiosyncratic shock $\varepsilon_{i,t}$ are IID standard normal. Then time-variation in the shape of $E_t[M_{t+1}|R_{t+1}]$ can be characterized by

$$\frac{\partial \ln E_t[M_{t+1}|R_{t+1}]}{\partial \ln R_{t+1}} = -\frac{\sigma_{m,t}}{\sigma_{s,t}} \times \frac{\sigma_{s,t}^2}{\sigma_{s,t}^2 + \sigma_{i,t}^2}. \tag{11}$$

Proof. *Joint normality implies that the log pricing kernel, conditional on the log return, is distributed as* $\ln M_{t+1}|\ln R_{t+1} \sim N\left(E_t[\ln M_{t+1}|\ln R_{t+1}], \mathrm{Var}_t[\ln M_{t+1}|\ln R_{t+1}]\right)$, *where* $E_t[\ln M_{t+1}|\ln R_{t+1}] = \mu_{m,t} - \frac{\sigma_{m,t}\sigma_{s,t}}{\sigma_{s,t}^2+\sigma_{i,t}^2}\left(\ln R_{t+1} - \mu_r\right)$ *and* $\mathrm{Var}_t[\ln M_{t+1}|\ln R_{t+1}] = \sigma_{m,t}^2\left(1 - \frac{\sigma_{s,t}^2}{\sigma_{s,t}^2+\sigma_{i,t}^2}\right)$. *Using the moment generating function of a normal random variable, the conditional expectation of the pricing kernel equals*

$$\mathbb{E}_t[M_{t+1}|R_{t+1}] = \exp\left\{\mu_{m,t} - \frac{\sigma_{m,t}\sigma_{s,t}}{\sigma_{s,t}^2 + \sigma_{i,t}^2}\left(\ln R_{t+1} - \mu_r\right) + \sigma_{m,t}^2\left(1 - \frac{\sigma_{s,t}^2}{\sigma_{s,t}^2 + \sigma_{i,t}^2}\right)/2\right\}.$$

Taking logs and differentiating with respect to $\ln R_{t+1}$ *yields* (11). ■

In Proposition 1 and our description of structural models below, we refer to a return shock as *idiosyncratic* if it is independent of the pricing kernel and as *systematic* otherwise. The variance of returns reflects both types of shocks and is given by $\mathrm{Var}_t[\ln R_{t+1}] = \sigma_{i,t}^2 + \sigma_{s,t}^2$. In order to find $\sigma_{s,t}^2$ in a model where $\ln M_{t+1}$ and $\ln R_{t+1}$ share multiple sources of risk, such as Bansal and Yaron (2004), we first need to map the model to the dynamics in Proposition 1. This can be done by linearly projecting $\ln R_{t+1}$ on $\ln M_{t+1}$ and computing the part of return variance that is attributable to the pricing kernel. Specifically, systematic return variance equals the product of the squared regression slope coefficient, $(cov_t[\ln R_{t+1}, \ln M_{t+1}]/\sigma_{m,t}^2)^2$, and the variance of the independent variable, $\sigma_{m,t}^2$, i.e., it is given by $\sigma_{s,t}^2 = cov_t[\ln R_{t+1}, \ln M_{t+1}]^2/\sigma_{m,t}^2$.

To understand the intuition behind Proposition 1, note that $-\frac{\sigma_{m,t}}{\sigma_{s,t}}$ equals the slope of $\ln E_t[M_{t+1}|R_{t+1}]$ in a world without idiosyncratic risk ($\sigma_{i,t} = 0$). In this world, log returns and the log pricing kernel are perfectly correlated and all else equal, an increase in return volatility (due to $\sigma_{s,t}$) makes the projection flatter. The intuitive reason is that a drop in the market is more indicative of high marginal

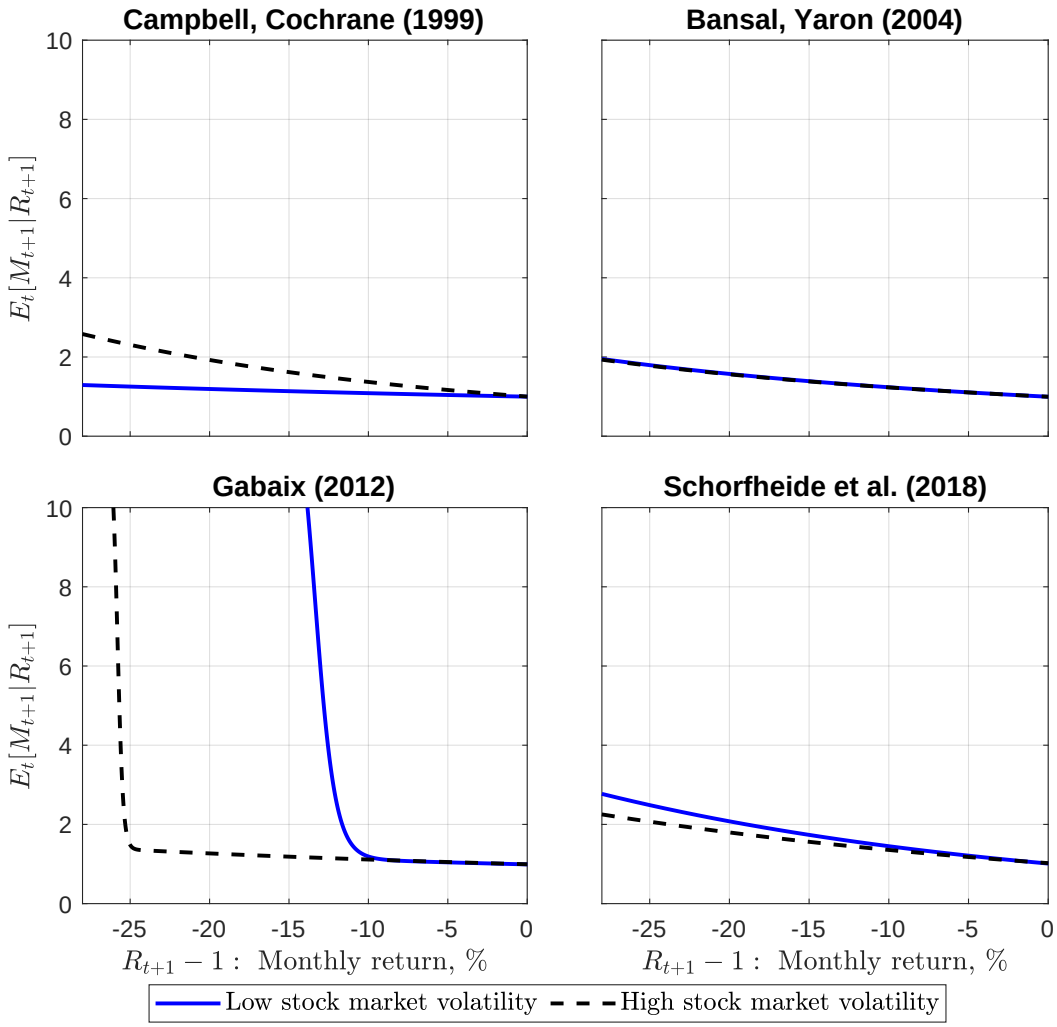

Figure VI: $E[M|R]$ in macrofinance models. We plot the projected pricing kernel for the 10th and 90th percentile of the conditional volatility of returns. The models of Campbell and Cochrane (1999), Bansal and Yaron (2004), and Schorfheide et al. (2018) are calibrated as in the original studies, whereas the model of Gabaix (2012) is calibrated as in Dew-Becker et al. (2017).

utility (deteriorating macroeconomic conditions) if it occurs during calm markets, rather than the middle of a recession.[14] The second factor in (11), $\frac{\sigma_{s,t}^2}{\sigma_{s,t}^2+\sigma_{i,t}^2}$, equals the fraction of systematic return risk or, equivalently, the squared conditional correlation between $\ln M_{t+1}$ and $\ln R_{t+1}$. Intuitively, an increase in the fraction of systematic return risk makes returns more informative about the pricing kernel and therefore results in a steeper projection.

Proposition 1 shows that time-variation in the shape of $E_t[M_{t+1}|R_{t+1}]$ reflects the joint evolution of $(\sigma_{m,t}, \sigma_{s,t}, \sigma_{i,t})$ over time. These dynamics could take many forms, but several stylized cases are noteworthy. First, suppose that idiosyncratic risk is time-invariant (constant $\sigma_{i,t}$), while systematic risk increases less than proportionally with $\sigma_{m,t}$. Then an increase in return volatility (due to $\sigma_{s,t}$, since $\sigma_{i,t}$ is constant) increases both fractions in equation (11) and therefore makes $E_t[M_{t+1}|R_{t+1}]$ steeper. We show below that this case describes the habit model. Second, suppose that $\sigma_{m,t}$, $\sigma_{s,t}$, and $\sigma_{i,t}$ are all proportional to one another. Then both fractions in equation (11) are time-invariant and the shape of $E_t[M_{t+1}|R_{t+1}]$ remains unchanged when return volatility rises. We show below that this case closely approximates the long-run risks model. Third, suppose that the volatility of returns evolves independently from $\sigma_{m,t}$. Then an increase in return volatility due to $\sigma_{i,t}$ will always make $E_t[M_{t+1}|R_{t+1}]$ flatter, whereas an increase due to $\sigma_{s,t}$ will make $E_t[M_{t+1}|R_{t+1}]$ flatter as long as $\sigma_{s,t} > \sigma_{i,t}$. Even though the model of Gabaix (2012) is not log-normal, we show below that its implications for time-variation in $E_t[M_{t+1}|R_{t+1}]$ are well-described by this scenario. A forth case that approximates case three occurs when $\sigma_{m,t} = a_0 + a_1 S_t$ and $\sigma_{s,t} = b_0 + b_1 S_t$ are linear functions of a state vector S_t with $a_0/b_0 > a_1/b_1$. This condition implies that $\sigma_{s,t}$ increases more than proportionally in $\sigma_{m,t}$. As a result, changes in return volatility are associated with smaller

[14] As an example, suppose that a -10% return corresponds to a 2 standard deviation event in normal times and, as such, tends to coincide with a 2 standard deviation shock to $M_{t=1}$ – a large macroeconomic contraction or significant news about the economic state. If return volatility doubles in recessions while pricing kernel volatility remains unchanged, then a return of -10% corresponds to a 1 standard deviation event and, as such, tends to coincide with a 1 standard deviation shock in M_{t+1} – a smaller macroeconomic contraction or less significant news about the economic state.

changes in pricing kernel volatility, and the first fraction in (11) decreases when return volatility rises. We show below that this case arises due to time preference shocks in the model of Schorfheide et al. (2018).

In what follows, we briefly review each of the aforementioned models and discuss what their economic mechanisms imply about the joint dynamics of $(\sigma_{m,t}, \sigma_{s,t}, \sigma_{i,t})$ and the resulting variation in the shape of $E_t[M_{t+1}|R_{t+1}]$.

B. Habits

B.1. Campbell and Cochrane (1999): Model statement

Aggregate consumption and dividends follow homoscedastic random walks,

$$\begin{aligned} \Delta c_{t+1} =& g + \sigma \varepsilon_{c,t+1} \\ \Delta d_{t+1} =& g + \sigma_w \left(\rho \varepsilon_{c,t+1} + \sqrt{1-\rho^2} \varepsilon_{d,t+1} \right), \end{aligned} \tag{12}$$

where ε_c and ε_d are IID standard normal. Equation (12) implies that the correlation between consumption and dividend growth rates equals ρ.[15] Equity represents a claim to the dividends in all future periods. The representative agent's utility function is

$$E_t \left[\sum_{h=0}^{\infty} \delta^h \frac{(C_{t+h} - X_{t+h})^{1-\gamma} - 1}{1-\gamma} \right], \tag{13}$$

where $\delta > 0$ controls time preference and $\gamma > 0$ controls risk preference. Time variation in the habit, X_t, is modelled via the surplus consumption ratio

$$S_t = \frac{C_t - X_t}{C_t}, \tag{14}$$

whose logarithm, $s_t = \ln S_t$, evolves via the heteroscedastic AR(1) process,

$$\begin{aligned} s_{t+1} =& (1-\phi)\bar{s} + \phi s_t + \lambda(s_t)\sigma \varepsilon_{c,t+1} \\ \lambda(s_t) =& \begin{cases} \frac{1}{\bar{S}} \sqrt{1 - 2(s_t - \bar{s})} - 1 & , s_t < s_{max} \\ 0 & , s_t \geq s_{max} \end{cases}, \end{aligned} \tag{15}$$

[15] We express dividend growth as a function of the consumption shock and an independent normal shock, $\varepsilon_{d,t+1}$. This is mathematically identical to the formulation in Campbell and Cochrane (1999), which relies on a single shock to dividends that is correlated with consumption, but makes it easier to distinguish between the idiosyncratic and systematic components of dividend risk.

where $\bar{S} = \sigma\sqrt{\gamma/(1-\phi)}$ and $s_{max} = \bar{s} + (1-\bar{S}^2)/2$. The parameters $\bar{s} = \ln \bar{S}$ and ϕ control the unconditional mean and persistence of s_t, whereas the function $\lambda(s_t)$ controls its sensitivity to consumption innovations. The model's implications for stock market risk are reflected in the price-dividend (P/D) ratio, which is an endogenous function of the model's state variable s_t.

B.2. Volatility Dynamics and $E[M|R]$ in the Habit Model

The conditional variance of the log pricing kernel, $\ln M_{t+1} = \ln \delta - \gamma(\Delta c_{t+1} + \Delta s_{t+1})$, is a linearly decreasing function of s_t,

$$\sigma^2_{m,t} = \gamma(1-\phi)(1+2\bar{s}) - 2\gamma(1-\phi)s_t, \tag{16}$$

whereas the conditional variance of idiosyncratic shocks $(\varepsilon_{d,t+1})$ is time-invariant,

$$\sigma^2_{i,t} = \sigma^2_w(1-\rho^2) \tag{17}$$

Systematic shocks $(\varepsilon_{c,t+1})$ are reflected in dividend growth rates, as well as time-variation in the P/D ratio. Because P/D has to be computed numerically, its conditional variance cannot be expressed analytically. Schreindorfer (2023) shows that, for a numerically accurate solution, the conditional variance of P/D looks drastically different from its counterpart in Campbell and Cochrane (1999). We first explain the model mechanism for the original (inaccurate) solution, which readers are more likely to be familiar with, and then point out what changes when the model is solved accurately. The projection depicted in Figure VI reflects the accurate solution. Note that any change in the volatility of returns, $\sqrt{\sigma^2_{i,t} + \sigma^2_{s,t}}$, reflects a change in $\sigma_{s,t}$ because $\sigma_{i,t}$ is time-invariant. The dashed line in the top-left panel of Figure VII shows that, for the original (inaccurate) solution, the conditional variance of systematic shocks, $\sigma^2_{s,t}$, is a concave and monotonically increasing function of $\sigma^2_{m,t}$, i.e., it increases less than proportionally with $\sigma^2_{m,t}$. Proposition 1 implies that, as a result, a rise in the volatility of returns makes $E[M|R]$ steeper, rather than flatter, as in the data.

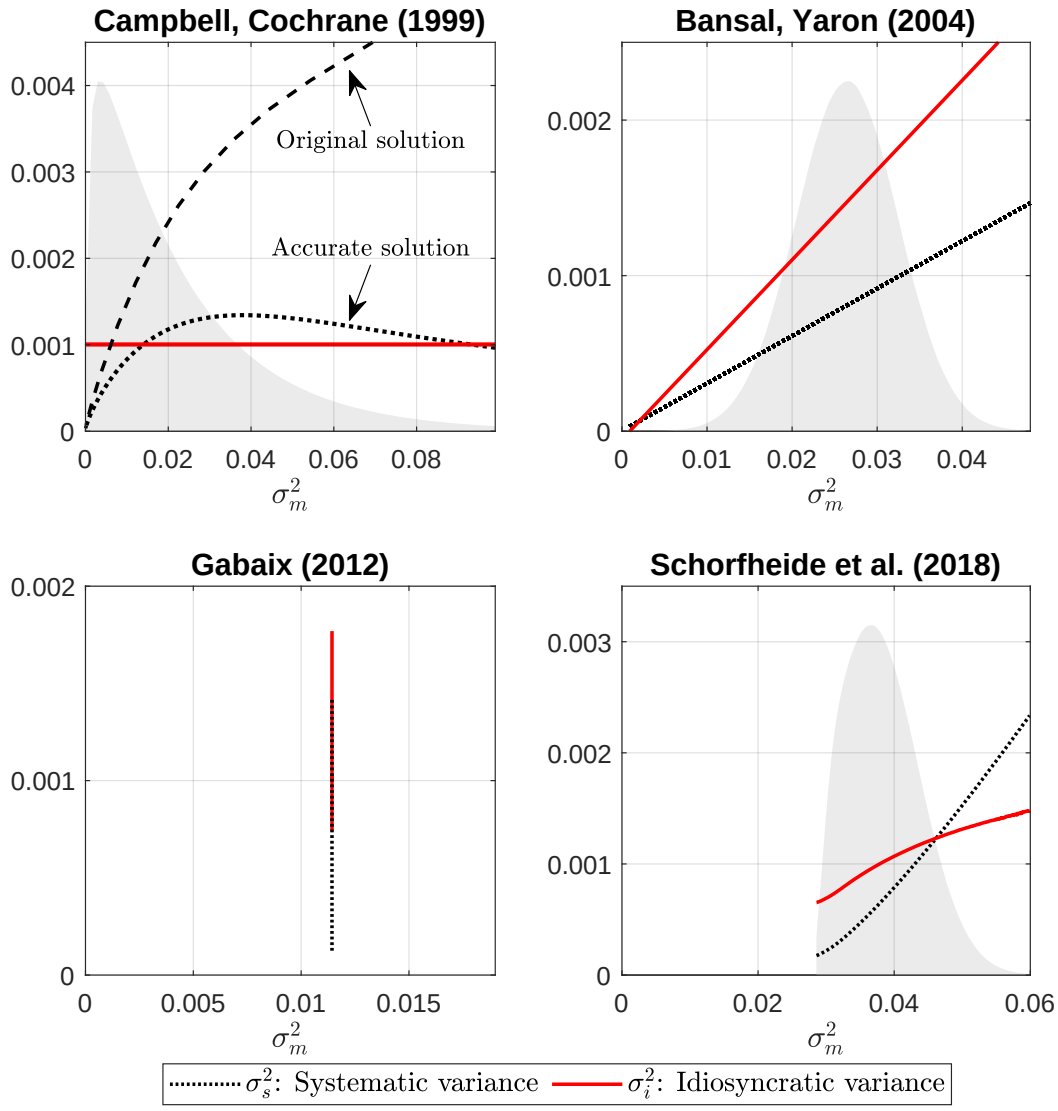

Figure VII: Sources of time-variation in $E[M|R]$. We plot the conditional variances of systematic and idiosyncratic return shocks as functions of the pricing kernel's conditional variance. Proposition 1 connects these moments to the shape of $E[M|R]$. The models of Campbell and Cochrane (1999), Bansal and Yaron (2004), and Schorfheide et al. (2018) are calibrated as in the original studies, whereas the model of Gabaix (2012) is calibrated as in Dew-Becker et al. (2017).

To see why these volatility dynamics are essential for the habit mechanism, note that a sequence of adverse consumption shocks leads to a decrease in s_t and an increase in $\sigma_{m,t}$. This raises risk aversion and risk premia and lowers the P/D ratio, resulting in a negative realized return. The simultaneous increase in $\sigma_{s,t}$ makes stocks riskier and thereby further increases risk premia. As a result, the model captures (i) the countercyclical nature of stock market volatility, i.e., the fact that volatility increases after adverse macroeconomic shocks (negative values of ε_c) (ii) the "leverage effect", i.e., the negative correlation between changes in volatility and contemporaneous returns (Black 1976), and (iii) the long-horizon predictability of excess returns, i.e., the observation that low values of P/D coincide with high risk premia. The habit model's counterfactual implications for time-variation in $E[M|R]$ therefore result directly from its core mechanism for dynamic asset pricing facts.

For the accurate solution, $\sigma^2_{s,t}$ is a hump-shaped (rather than monotonically increasing) function of $\sigma^2_{m,t}$, as shown in Figure VII. This alters the model's predictions in several ways relative to the original solution. First, during recessions (states of high $\sigma^2_{m,t}$), a further increase in $\sigma_{m,t}$ decreases $\sigma_{s,t}$ and therefore decreases the conditional volatility of returns, making stocks safer. When solved accurately, the habit model is therefore no longer consistent with (i) the fact that stock market volatility spikes during recessions, (ii) the leverage effect, or (iii) the strong predictability of long-horizon returns (see Schreindorfer 2023 for more details). Second, states of low return volatility (low $\sigma_{s,t}$) can be associated with either low or high values of $\sigma_{m,t}$ for the accurate solution, and hence with either a flat or steep $E[M|R]$ curve. Probabilistically, however, values around the 10th percentile of return volatility are more than 100,000 times as likely to coincide with low values of σ_m and a flat $E[M|R]$ curve than with high values of σ_m and a steep $E[M|R]$ curve, whereas values around the 90th percentile of return volatility always coincide with a steep $E[M|R]$ curve. Even for the accurate solution, for which we depict $E[M|R]$ in Figure VI, the habit mechanism therefore remains inconsistent with our empirical evidence.

C. Recursive Utility and Long-run Risks

C.1. Bansal and Yaron (2004): Model statement

The representative agent has Epstein-Zin utility, calibrated to imply a preference for the early resolution of uncertainty (EIS>1/RRA), and is endowed with

$$\begin{aligned}
\Delta c_{t+1} =& g + x_t + \sigma_{c,t}\eta_{t+1} \\
\Delta d_{t+1} =& g + \phi x_t + \varphi_d \sigma_{c,t} u_{t+1} \\
x_{t+1} =& \rho x_t + \varphi_e \sigma_{c,t} e_{t+1} \\
\sigma^2_{c,t+1} =& \sigma^2 + \nu_1(\sigma^2_{c,t} - \sigma^2) + \sigma_w w_{t+1},
\end{aligned}$$

where $(\eta_{t+1}, u_{t+1}, e_{t+1}, w_{t+1}) \overset{\text{IID}}{\sim} N(0,1)$. The model's log-linearized solution implies that the log price-dividend ratio, $\ln(P_t/D_t) = A_{0m} + A_{1m}x_t + A_{2m}\sigma^2_{c,t}$, is a function of the state variables x_t and $\sigma_{c,t}$, which generate persistent variation ("long run risks") in the conditional mean and volatility of consumption and dividend growth rates. The log pricing kernel,

$$\ln M_{t+1} = E_t[\ln M_{t+1}] + \lambda_{m,\eta}\sigma_{c,t}\eta_{t+1} - \lambda_{m,e}\sigma_{c,t}e_{t+1} - \lambda_{m,w}\sigma_w w_{t+1}, \tag{18}$$

and the log ex-dividend return,

$$\ln R_{t+1} = E_t[\ln R_{t+1}] + A_{1m}\varphi_e\sigma_{c,t}e_{t+1} + A_{2m}\sigma_w w_{t+1} + \varphi\sigma_{c,t}u_{t+1}, \tag{19}$$

are linear functions of normal shocks, and therefore conditionally jointly log-normal. We refer readers to Bansal and Yaron (2004) for expressions of the endogenous coefficients $(\lambda_{m,\eta}, \lambda_{m,e}, \lambda_{m,w}, A_{0m}, A_{1m}, A_{2m})$.

C.2. Volatility Dynamics and $E[M|R]$ in the Long-run Risks Model

The conditional variances of the log pricing kernel and log return follow from (18) and (19) as

$$\begin{aligned}
\sigma^2_{m,t} =& \lambda^2_{m,\eta}\sigma^2_{c,t} + \lambda^2_{m,e}\sigma^2_{c,t} + \lambda^2_{m,w}\sigma^2_w \\
\sigma^2_{r,t} =& A^2_{1m}\varphi^2_e\sigma^2_{c,t} + A^2_{2m}\sigma^2_w + \varphi^2\sigma^2_{c,t}.
\end{aligned}$$

Since there are two systematic return shocks (e_{t+1} and w_{t+1}), we compute the variance of systematic returns as the squared ratio of $cov_t[\ln R_{t+1}, \ln M_{t+1}]$ and $\sigma_{m,t}$,

$$\sigma_{s,t}^2 = (-\lambda_{m,e} A_{1m} \varphi_e \sigma_{c,t}^2 - \lambda_{m,w} A_{2m} \sigma_w^2)^2 / \sigma_{m,t}^2.$$

The variance of idiosyncratic returns, $\sigma_{i,t}^2$, equals the difference between $\sigma_{r,t}^2$ and $\sigma_{s,t}^2$. These expressions show that time-variation in the model's state variable $\sigma_{c,t}$ is the only source of variation in the conditional volatility of returns and the pricing kernel. For Bansal and Yaron's calibration, $\sigma_{m,t}^2$, $\sigma_{s,t}^2$, and $\sigma_{i,t}^2$ are approximately proportional to one another, as shown in the top-right panel of Figure VII. This proportionality explains, via Proposition 1, why the shape of $E[M|R]$ is time-invariant in the model.

The positive relation between $\sigma_{m,t}^2$ and $\sigma_{s,t}^2$ is central to the long-run risks mechanism. As in the habit model, it implies that elevated stock market volatility coincides with a large equity premium and a low P/D ratio. The long-run risks model therefore explains the negative correlation between volatility innovations and contemporaneous returns, as well as the long-horizon predictability of excess returns. Additionally, because Epstein-Zin preferences imply aversion against heteroscedasticity in x_t, a positive realization of $\varepsilon_{\sigma,t}$ increases the pricing kernel and lowers P/D (contributing to the equity premium), while also increasing $\sigma_{m,t}^2$ and $\sigma_{s,t}^2$. Covariation between $\sigma_{m,t}^2$ and $\sigma_{s,t}^2$ is therefore a byproduct of how the model generates the equity premium. Bansal and Yaron (2004, Table 5) show that about 40% of the pricing kernel's variance, and hence about 40% of the equity premium, is attributable to the time-varying volatility channel. Positive comovement between $\sigma_{m,t}^2$ and $\sigma_{s,t}^2$ is therefore essential for the long-run risks model's main empirical predictions.

The fact that $\sigma_{i,t}^2$ covaries with $\sigma_{m,t}^2$ and $\sigma_{s,t}^2$ is not important, because the long-run risks factor $\sigma_{c,t}^2$ affects P/D primarily by controlling the conditional volatility of x_t, not by controlling the conditional volatility of consumption and dividend growth rates. If consumption and dividend growth rates were homoscedastic, $\sigma_{i,t}^2$ would be time-invariant and $E[M|R]$ would behave as in the Campbell-Cochrane model, but

the model's asset pricing implications would remain nearly unchanged.[16] Obviously, this change would remove the model further from explaining our empirical evidence.

C.3. Related Recursive Utility Models

The online appendix IA.III examines a number of related models with recursive preferences and persistent state variables, some of which were explicitly designed to capture features of index option markets. First, we consider Drechsler and Yaron (2011), who augment the long-run risks model with jumps in x_t and $\sigma^2_{c,t}$ and a second volatility factor. These changes result in left-skewed returns and a better fit to volatility moments, but do not alter the model's main economic mechanism. We find that they also have little effect on $E[M|R]$, which remains close to time-invariant.

Second, we consider Wachter (2013), who augments the disaster model of Rietz (1988) and Barro (2006) with recursive utility and persistent variation in the disaster probability, which is modelled as a Cox et al. (1985) process. The CIR process implies that a rise in the disasters probability not only lowers expected consumption growth, but also increases the conditional volatility of expected consumption growth. Hence, Wachter essentially fuses Bansal and Yaron's x_t and $\sigma^2_{c,t}$ into a singe state variable. As in the original long-run risks model, $\sigma^2_{m,t}$ and $\sigma^2_{s,t}$ are approximately proportional to one another and $E[M|R]$ is time-invariant as a result.

Lastly, we consider Constantinides and Ghosh (2017), who propose a quantitative implementation of the Constantinides and Duffie (1996) mechanism. Their model features recursive preferences and a persistent state that simultaneously controls time-variation in (i) the higher moments of households' idiosyncratic income shocks and (ii) the conditional mean of aggregate dividend growth. The state is modelled

[16] For confirmation, we solved the model with homoscedastic growth rates by replacing $\sigma_{c,t}$ by σ (the square root of the unconditional mean of $\sigma^2_{c,t}$) in the equations for Δc_{t+1} and Δd_{t+1}, while maintaining the time-varying $\sigma_{c,t}$ in the equation for x_{t+1}. Based on a long simulation, the mean/std/AC1 of the annual (time-aggregated) log P/D ratio equal 3.04/0.18/0.69 for both the heteroscedastic and homoscedastic versions of the model, whereas the annual log equity return has a mean/std/AC1 of 6.68/16.62/0.02 in the heteroscedastic model and 6.69/16.73/0.02 in the homoscedastic model.

as an autoregressive Gamma process, which implies, similar to the CIR process in Wachter (2013), that the state's volatility increases in its level. When household risk rises and increases $\sigma^2_{m,t}$, the conditional mean of dividends simultaneously becomes more volatile, which increases the volatility of P/D and therefore $\sigma^2_{s,t}$. Once again, this mechanism implies that $\sigma^2_{m,t}$ and $\sigma^2_{s,t}$ are approximately proportional to one another and that the $E[M|R]$ curve is time-invariant as a result.

These examples illustrate that it is the long-run risks mechanism for time-varying volatility, rather than the details of its implementation in Bansal and Yaron (2004), that lead to the models' inconsistency with the observed variation in $E[M|R]$.

D. Time-varying Disaster Resilience

D.1. Gabaix (2012): Model statement

We consider the version of Gabaix's (2012) time-varying disaster risk model in Dew-Becker et al. (2017).[17] The representative agent has time-separable power utility and is endowed with

$$\begin{aligned} \Delta c_{t+1} &= g + \sigma\varepsilon_{t+1} + \xi_{t+1}J_{t+1} \\ \Delta d_{t+1} &= g + \lambda\sigma\varepsilon_{t+1} - L_{t+1}\xi_{t+1} \end{aligned} \tag{20}$$

where $\varepsilon_{t+1} \overset{\text{IID}}{\sim} N(0,1)$ captures small consumption shocks, $\xi_{t+1} \overset{\text{IID}}{\sim} \text{Bernoulli}(p_J)$ is a disaster indicator, and $J_{t+1} \overset{\text{IID}}{\sim} N(E_J, V_J^2)$ is the disaster size. The recovery rate of dividends, L_t, follows the autoregressive process

$$L_{t+1} = (1-\rho_L)\bar{L} + \rho_L L_t + \sigma^L \varepsilon_{L,t+1}, \tag{21}$$

where $\varepsilon_{L,t+1} \overset{\text{IID}}{\sim} N(0,1)$. Consumption disasters lead to a large drop in dividends if they occur in periods with high L_t and a small drop in dividends if they occur in periods with low L_t. Different from the model of Wachter (2013), but as in the

[17] Dew-Becker et al.'s version of the model adds Gaussian innovations to consumption growth (which is constant absent disasters in the model's original version), it assumes a normally distributed (rather than constant) disaster size, and it specifies the recovery rate of dividends as an autoregressive (rather than linearity generating) process. These modifications do not alter the model's basic economics, but result in a pricing kernel that is a continuous (rather than discontinuous) function of the model's state. We follow Dew-Becker et al.'s calibration of the model.

original disaster model of Rietz (1988) and Barro (2006), the probability of disasters, p_J, is assumed to be constant.

D.2. Volatility Dynamics and $E[M|R]$ in the Rare Disaster Model

The bottom-left panel of Figure VI shows that Gabaix's model not only captures $E[M|R]$'s time-variation, but also its steep slope in the left tail of the return distribution. These implications reflect four forces. First, consumption disasters are associated with a very high value of the pricing kernel M_{t+1}, whereas non-disaster states are not. Second, dividends and returns drop by $-L_{t+1}$ when a disaster materializes. If disasters were the only shock affecting dividends, returns of $-L_{t+1}$ would always coincide with high marginal utility and returns of 0 with low marginal utility. Third, the small Gaussian shocks ε_{t+1} and $\varepsilon_{L,t+1}$ add "noise" to dividends and returns. In their presence, returns to the far left of $E_t[-L_{t+1}]$ are almost sure to coincide with a disaster realization a high value of M_{t+1}, returns to the far right of $E_t[-L_{t+1}]$ are almost sure to coincide with no disaster realization a low value of M_{t+1}, and returns in the vicinity of $E_t[-L_{t+1}]$ could coincide with either. The Gaussian shocks have a small conditional volatility, so that the increase in $E_t[M_{t+1}|R_{t+1}]$ from low to high marginal utility states around $E_t[-L_{t+1}]$ occurs fairly rapidly. Lastly, because $E_t[L_{t+1}]$ changes over time, so does the return region for which the pricing kernel increases steeply. Specifically, periods of high L_t (and high $E_t[-L_{t+1}]$) are associated with higher conditional return volatility, as well as larger negative returns in case of a disaster. The steep increase in $E[M|R]$ therefore occurs for more extreme negative return outcomes when volatility is low.

The mechanism can alternatively be cast in the language of Proposition 1.[18] For that purposes, the most noteworthy feature of Gabaix's model is that the volatility of returns evolves independently from all shocks that affect consumption or the pricing

[18]The variance of systematic return shocks has to be computed numerically because the P/D ratio is not available in closed form. We compute $\sigma^2_{s,t}$ based on the ratio of $cov_t[\ln R_{t+1}, \ln M_{t+1}]^2$ and $\sigma^2_{m,t}$, compute the variance of idiosyncratic return shocks as $\sigma^2_{i,t} = \sigma^2_{r,t} - \sigma^2_{s,t}$, and plot both in the bottom-left panel of Figure VII. Because Gabaix's model is not lognormal, this decomposition into idiosyncratic and systematic return risk only serves as an approximation in this case.

kernel. This property implies that changes in return volatility are not accompanied by changes in $\sigma_{m,t}$. Proposition 1 shows that, under lognormality, an increase in return volatility that is not accompanied by a change in $\sigma_{m,t}$ makes $E[M|R]$ flatter, as in the data. Despite the fact that Gabaix' model is not lognormal, Proposition 1 describes its qualitatively implications for time-variation in $E[M|R]$.

Two additional points are worth making about Gabaix (2012). First, the fact that M_{t+1} is IID is not important for the model's ability to rationalize our evidence. What matters is that the pricing kernel's conditional variance $\sigma^2_{m,t}$ does not comove with the conditional variance of returns, as one can see from Proposition 1. Second, what matters for the model's ability to explain our evidence is not that shocks to the variance of returns are independent from M_{t+1}, but rather that they are independent from shocks to the variance of M_{t+1}. This is a subtle but important distinction. It is conceivable that shocks to return variance affect the realized value of the pricing kernel without altering its conditional variance. This would rationalize our finding and Jurado et al.'s (2015) observation that macroeconomic uncertainty (a plausible driver of pricing kernel volatility) is only weakly correlated with return volatility, but not the fact that shocks to conditional return volatility are neither priced (Dew-Becker et al. 2017) nor predictive of future economic contractions (Berger et al. 2020). Assuming that the variance of returns is independent of both M_{t+1} and its conditional moments, as in Gabaix's model, is an intuitive joint explanation for all of these findings.

E. Valuation Risk

E.1. Time Preference Shocks

Albuquerque et al. (2016) augment the recursive preference specification of Epstein and Zin (1989) with stochastic variation in the rate of time preference to capture "valuation risk" due to demand shocks. Utility is given by

$$V_t = \left[(1-\delta)\lambda_t C_t^{\frac{1-\gamma}{\theta}} + \delta \left(\mathbb{E}\left[V_{t+1}^{1-\gamma} \right] \right)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \tag{22}$$

where λ_t determines how agents trade off current versus future utility. The log growth rate of the preference shifter, $x_{\lambda,t+1} = \ln\left(\frac{\lambda_{t+1}}{\lambda_t}\right)$, evolves as an AR(1) process with shocks that are independent of cash flow shocks,[19]

$$x_{\lambda,t+1} = \rho_\lambda x_{\lambda,t} + \sigma_\lambda \varepsilon_{\lambda,t+1}. \tag{23}$$

Albuquerque et al. (2016) show that the augmented preference specification can account for the weak correlation between stock returns and fundamentals and generate realistic risk-free rate dynamics. In the following, we illustrate how time preference shocks affect the projected pricing kernel in the model of Schorfheide et al. (2018), who embed them into an economy with long-run risks.[20]

E.2. Volatility Dynamics and $E[M|R]$ in the Model of Schorfheide et al. (2018)

Schorfheide et al. (2018) augment the original long-run risks model with the aforementioned preference shocks, non-zero correlation between consumption and dividend innovations, and separate volatility processes for consumption, dividends, and the conditional mean of growth rates. We refer interested readers to online Appendix IA.IV for a more detailed discussion of their model.

The bottom right panel of Figure VI shows that the model produces time-variation in $E[M|R]$ that is qualitatively in line with the data. To illustrate the source of this time-variation, we solve the model without a preference shock (but otherwise identical parameters) by setting its volatility σ_λ to zero. The left panel of Figure VIII shows that this alternative calibration generates almost no time-variation in the projected kernel, similar to the original long-run risks model. For our purposes, preference shocks therefore represent the model's key deviation from Bansal and Yaron (2004).

[19] We rely on the notation and timing convention of Schorfheide et al. (2018), rather than those of Albuquerque et al. (2016).

[20] In unreported results, we examine $E[M|R]$ in the model of Albuquerque et al. (2016). While preference shocks have a similar qualitative effect in their model as in Schorfheide et al. (2018), the effect is quantitatively much smaller because the mode of Albuquerque et al. (2016) generates almost no heteroscedasticity in returns.

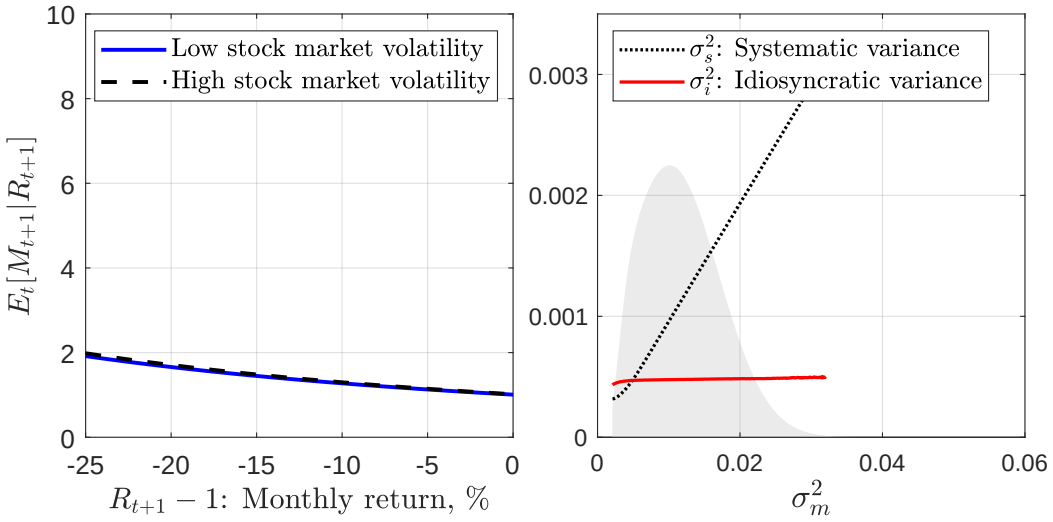

Figure VIII: Schorfheide et al. (2018) without preference shock. We plot the analogue of to Figure VI and Figure VII, respectively, for the Schorfheide et al. (2018) model where we eliminate the preference shock by setting $\sigma_\lambda = 0$.

Because the model is conditionally log-normal, we can rely on Proposition 1 to understand its implications for time-variation in $E[M|R]$. The bottom-right panel of Figure VII shows σ_s^2 and σ_i^2 as functions of σ_m^2 for the calibration of Schorfheide et al. (2018), whereas the right panel of Figure VII shows the same plot for our calibration without preference shocks. The main deviation from Bansal and Yaron (2004) is immediately obvious: The conditional variance of the pricing kernel, σ_m^2 never falls below a value of 0.028 in the calibration with preference shocks, i.e., it is of the form $\sigma_m^2 = a_0 + a_1 S_t$ with a large constant term and moderate dependence on the vector of state vector S_t. The large constant term implies that σ_s^2 increases more than proportionally in σ_m^2, which leads to a flattening of $E[M|R]$ via the first term in Proposition 1. In contrast, the calibration without preference shocks implies that the constant term in σ_m^2 is close to zero, as in the original long-run risks model. It is worth noting that the large constant term in σ_m^2 can be interpreted as an approximation to the case in which the volatility of returns is independent from that of the pricing kernel, because it implies that relative changes in σ_m^2 are small compared to relative changes in σ_r^2.

IV. Conclusion

Option markets provide us with invaluable information about the pricing of stock market risks. We show that negative returns are substantially more painful to investors when they occur in periods of low volatility, which is reflected in a steeper projected pricing kernel and larger risk premia on out-of-the-money put options. Our finding provides a useful diagnostic test for asset pricing models, which routinely assume difficult-to-measure dynamics in preferences and fundamentals to rationalize asset prices. We show that models with habits or recursive utility require counterfactual pricing kernel behavior to explain asset price dynamics, such as countercyclical volatility, the leverage effect, or long-horizon return predictability.

We trace the source of this discrepancy back to the joint dynamics of returns and pricing kernel. Our Proposition 1 shows that return volatility must evolve (close to) independently of the volatility of the pricing kernel to explain the observed variation in $E[M|R]$. This theoretical result aligns with prior empirical evidence. First, Jurado et al. (2015) show that macroeconomic uncertainty (a plausible driver of pricing kernel volatility) is only weakly correlated with return volatility. Second, Glosten et al. (1993) find that expected returns (which are positively related to pricing kernel volatility) do not comove significantly with stock market volatility.

Our findings suggest that, in order to capture the projected pricing kernel's time-variation quantitatively, it is necessary to simultaneously capture its strong convexity. The model of Gabaix (2012) stands apart from other mainstream theories in its ability to capture both features. To do so, it combines the rare disaster mechanism with a source of return heteroscedasticity (time-varying disaster exposure) that evolves independent of the pricing kernel. An interesting avenue for future work would be to explore whether a similar mechanism can be used in other asset pricing paradigms.

Appendix

This appendix explains how we extract risk-neutral distributions from option prices and details the time series model for conditional volatility.

A. Extracting Risk-neutral Densities from Options

We follow the methodology in Beason and Schreindorfer (2022) to extract risk-neutral densities from option prices. Breeden and Litzenberger (1978) show that the risk-neutral PDF of the future price level S_{t+1} is given by

$$f_t^*(S_{t+1}) = R_t^f \times \frac{\partial^2 P_t(K)}{\partial K^2}\bigg|_{K=S_{t+1}}, \tag{A.1}$$

where P is the price of a put option and K the associated strike price. The risk-neutral PDF of ex-dividend returns follows from the change of variables $R_{t+1} = \frac{S_{t+1}}{S_t}$ as $f_t^*(R_{t+1}) = S_t \times f_t^*(S_{t+1})$. To recover risk-neutral densities from options based on (A.1), it is necessary to observe option prices for the desired maturity and a continuum of strikes. We generate these prices via interpolation and extrapolation of observed quotes as follows. For each day in the sample, we use Black's formula (a version of Black and Scholes 1973) to convert observed option prices to implied volatility (IV) units, fit an interpolant to them, evaluate the interpolant at a maturity of 30 calendar days and a fine grid of strike prices, map interpolated IVs back to option prices, and finally compute f_t^* via finite differences based on (A.1). Importantly, this approach does not assume the validity of the Black-Scholes model because Black's formula is merely used to map back-and-forth between two spaces. The mapping relies on LIBOR rates that are linearly interpolated to the options' maturities and forward prices for the underlying. The remainder of this appendix details the interpolation of IVs.

The SVI Method

We interpolate IVs based on Jim Gatheral's SVI method.[21] SVI describes implied variance (the square of IV) for a given maturity τ with the function

$$\sigma^2_{BSM}(x) = a + b\left(\rho(x-m) + \sqrt{(x-m)^2+\sigma^2}\right), \tag{A.2}$$

where $x = \log(K/F_{t,\tau})$ is the option's log-moneyness, $F_{t,\tau}$ the forward price for maturity τ, and a, b, ρ, m, σ are parameters. The method is widely used in financial institutions because it is parsimonious, yet known to provide a very good approximation to IVs, both in the data and in fully-specified option pricing models.

We make two modifications to the basic SVI method to allow for interpolation in the maturity, in addition to the moneyness dimension. First, we parameterize σ^2_{BSM} as a function of standardized moneyness, $\kappa \equiv \frac{\log(K/F_{t,\tau})}{VIX_t/100\times\sqrt{\tau}}$, rather than x, to limit the extent to which the shape of the IV curve varies with maturity. Second, we specify linear functions of τ for the five coefficients, e.g.,

$$a = a_0 + a_1\tau, \tag{A.3}$$

and similarly for (b, ρ, m, σ). Jointly, (A.2) and (A.3) describe IVs as a bivariate function of κ and τ that is parameterized by $\theta \equiv (a_0, a_1, b_0, b_1, \rho_0, \rho_1, m_0, m_1, \sigma_0, \sigma_1)$.

An important criterion for the successful interpolation and extrapolation of IVs is that the corresponding option prices respect theoretical no arbitrage restrictions, i.e. that they are (i) non-negative, (ii) monotonic in K, (iii) convex in K, and (iv) imply (via Equation A.1) a density $f^*_t(R)$ that integrates to one. We impose these constraints in the estimation as further described below.

Data and Implementation

We clean the options data by removing observations that (i) violate the static no-arbitrage bounds $P \leq K/R^f$ or $C \leq S$, (ii) have a best bid quote of zero, (iii)

[21] SVI was devised at Merrill Lynch and disseminated publicly by Gatheral (2004). See Gatheral (2006) for a textbook treatment and Berger et al. (2020) for a recent application in economics.

have the CBOE's error code 999 for ask quotes or 998 for bid quotes, (iv) have non-positive bid-ask spreads, (v) have midquotes less than \$0.50, (vi) are singles (a call quote without a matching put quote or vice versa), (vii) are PM settled, or (viii) have IVs less than 2% or more than 200%. To detect any additional outliers, we fit a linear function κ and τ to IVs on each date, and remove observations that are highly influential based on their Cook's distance (a common statistical metric for detecting outliers). Finally, we restrict the sample to puts with a standardized moneyness below 0.5, calls with a standardized moneyness above -0.5, and maturities between 8 and 120 calendar days, i.e. we exclude long-term and in-the-money options.

For each day in the sample, we estimate the SVI parameter vector θ by minimizing the root mean squared error between observed IVs and the SVI interpolant,

$$\hat{\theta}_t = \underset{\theta}{\operatorname{argmin}} \sqrt{\frac{1}{N_t} \sum_{i=1}^{N_t} [\sigma_{BSM,t,i} - \sigma_{BSM}(\kappa_{t,i}, \tau_{t,i}; \theta)]^2}, \tag{A.4}$$

where N_t is the number of observations on day t. We use a particle swarm algorithm to minimize the objective function and discard parameters for which SVI-implied prices violate no arbitrage constraints. The positivity, monotonicity, and convexity of option prices are checked on a bivariate grid for κ and τ.[22] At every maturity in the τ-grid, we integrate f_t^* over the κ-region and discard parameters for which these integrals do not fall within (a numerical error tolerance of) 1 basis point of one.

The fit to IVs results in an average (median) R^2 of 98.8% (99.6%) across the 7,556 trading days in our sample.

B. Conditional Volatility Estimation

We estimate the conditional volatility of stock market returns based on the heterogeneous autoregressive (HAR) model of Corsi (2009):

$$RV_t^{(21)} = \alpha + \beta^m RV_{t-21}^{(21)} + \beta^w RV_{t-21}^{(5)} + \beta^d RV_{t-21}^{(1)} + \epsilon_t, \tag{B.1}$$

[22] The κ-grid includes the integers from -20 to -11, 61 equally-spaced points between -10 and 5, and the integers from 6 to 10, for a total of 76 points. The width of this grid ensures that even extrapolated option prices are arbitrage free. The τ-grid is equally-spaced with 12 points between 10 and 120 days to maturity.

where the realized volatility $RV_t^{(1)} = (\sum_{i=1}^{N_t} r_{it}^2)^{0.5}$ equals the square root of the sum of N_t squared intra-daily log returns on day t, and $RV_t^{(h)} = (\frac{1}{h}\sum_{j=0}^{h} RV_{t-j}^{(1)})^{0.5}$. The weekly ($RV_t^{(5)}$) and monthly ($RV_t^{(21)}$) volatility measures allow the model to capture the well-known long-memory feature of stock market volatility. As is common, we sample returns every five minutes to compute RV. Doing so provides a good trade-off between keeping noise low (which increases at low sampling frequencies) and also avoiding microstructure effects (which are more prevalent at high sampling frequencies). We sub-sample our estimator every minute, which reduces the noise without affecting the estimator's bias, and add the squared log overnight return to each intra-daily variance estimate.

Intra-daily returns are computed from future prices for the S&P 500 index, which we obtained from Tick Data Inc. In 1997, the Chicago Mercantile Exchange (CME) introduced the so-called mini future (symbol: ES). Over time, the standard "large" futures contract (symbol: SP) lost market share to the mini, and eventually was discontinued in 2021. Since the dollar trading volume of the mini overtook that of the large contract during 2002, we switch our RV calculation from the large contract to the mini in 2003.

Our volatility forecasts are out-of-sample and implemented based on an expanding estimation window. We start the sample in 1988, so that our first forecast on Jan 02, 1990 is based on a two year estimation window. The model forecasts volatility very well with an out-of-sample R^2_{OOS} of 60.4%.

References

Adrian, T., R. K. Crump, and E. Vogt (2019): "Nonlinearity and Flight-to-Safety in the Risk-Return Trade-Off for Stocks and Bonds," *Journal of Finance*, 74, 1931–1973.

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

Albuquerque, R., M. Eichenbaum, V. X. Luo, and S. Rebelo (2016): "Valuation risk and asset pricing," *The Journal of Finance*, 71, 2861–2904.

Aretz, K., M.-T. Lin, and S.-h. Poon (2022): "Moneyness, Volatility, and the Cross-Section of Option Returns," *Review of Finance, forthcoming*.

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

Barone-Adesi, G., R. F. Engle, and L. Mancini (2008): "A GARCH option pricing model with filtered historical simulation," *Review of Financial Studies*, 21, 1223–1258.

Barro, R. J. (2006): "Rare disasters and asset markets in the twentieth century," 121, 823–866.

Beason, T. and D. Schreindorfer (2022): "Dissecting the Equity Premium," *Journal of Political Economy*, 130, 2203–2222.

Berger, D., I. Dew-Becker, and S. Giglio (2020): "Uncertainty shocks as second-moment news shocks," *The Review of Economic Studies*, 87, 40–76.

Black, F. (1976): "Studies of stock price volatility changes," *Proceedings of the 1976 Meetings of the American Statistical Association*, 171–181.

Black, F. and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–654.

Bliss, R. R. and N. Panigirtzoglou (2004): "Option-implied risk aversion estimates," *Journal of Finance*, 59, 407–446.

Breeden, D. T. and R. H. Litzenberger (1978): "Prices of state-contingent claims implicit in option prices," *Journal of Business*, 51, 621–651.

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

Chapman, D. A. (1997): "Approximating the asset pricing kernel," *Journal of Finance*, 52, 1383–1410.

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, 228–265.

Christoffersen, P., S. Heston, and K. Jacobs (2013): "Capturing option anomalies with a variance-dependent pricing kernel," *Review of Financial Studies*, 26, 1963–2006.

Cochrane, J. H. (2005): *Asset pricing: Revised edition*, Princeton university press.

Constantinides, G. M. and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," 104, 219–240.

Constantinides, G. M. and A. Ghosh (2017): "Asset Pricing with Countercyclical Household Consumption Risk," *Journal of Finance*, 72, 415–460.

Corsi, F. (2009): "A simple approximate long-memory model of realized volatility," *Journal of Financial Econometrics*, 7, 174–196.

Cox, J. C., J. E. Ingersoll Jr, and S. A. Ross (1985): "A THEORY OF THE TERM STRUCTURE OF INTEREST RATES," *Econometrica (pre-1986)*, 53, 385.

Dew-Becker, I. and S. Giglio (2022): "Risk preferences implied by synthetic options," *Working Paper*.

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

Dittmar, R. F. (2002): "Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns," *Journal of Finance*, 57, 369–403.

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

Driessen, J., J. Koëter, and O. Wilms (2020): "Behavioral in the Short-run and Rational in the Long-run? Evidence from S&P 500 Options," *Working Paper, Tilburg University*.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Gabaix, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," 127, 645–700.

Gatheral, J. (2004): "A parsimonious arbitrage-free implied volatility parametrization with application to the valuation of volatility derivatives," Talk at Global Derivatives.

——— (2006): *The Volatility Surface: A Practitioner's Guide*, Wiley Finance.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993): "On the relation between the expected value and the volatility of the nominal excess return on stocks," *Journal of Finance*, 48, 1779–1801.

Hu, G. and K. Jacobs (2020): "Volatility and Expected Option Returns," *Journal of Financial and Quantitative Analysis*, 55, 1025–1060.

Jackwerth, J. C. (2000): "Recovering risk aversion from option prices and realized returns," *Review of Financial Studies*, 13, 433–451.

Jones, C. S. (2006): "A nonlinear factor analysis of S&P 500 index option returns," *Journal of Finance*, 61, 2325–2363.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Kim, H. J. (2022): "Characterizing the Conditional Pricing Kernel: A New Approach," *Working Paper, University of Houston*.

Linn, M., S. Shive, and T. Shumway (2018): "Pricing kernel monotonicity and conditional information," *Review of Financial Studies*, 31, 493–531.

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

Nagel, S. and K. J. Singleton (2011): "Estimation and evaluation of conditional asset pricing models," *The Journal of Finance*, 66, 873–909.

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

Rietz, T. A. (1988): "The equity risk premium: a solution," 22, 117–131.

Rosenberg, J. V. and R. F. Engle (2002): "Empirical pricing kernels," *Journal of Financial Economics*, 64, 341–372.

Schorfheide, F., D. Song, and A. Yaron (2018): "Identifying long-run risks: A Bayesian mixed-frequency approach," *Econometrica*, 86, 617–654.

Schreindorfer, D. (2023): "By Force of Habit and Cyclical Leverage," *Working Paper*.

Vuong, Q. H. (1989): "Likelihood ratio tests for model selection and non-nested hypotheses," *Econometrica: journal of the Econometric Society*, 307–333.

Wachter, J. A. (2013): "Can time-varying risk of rare disasters explain aggregate stock market volatility?" *Journal of Finance*, 68, 987–1035.