The Cyclicality of Risk and Risk Premia*

Christian Skov Jensen and Eben Lazarus

JULY 2025

Abstract

We consider and reexamine the properties of the market risk premium and variance over the business cycle. While it is well known that the risk premium and variance both increase in recessions, we find that the risk premium is less strongly countercyclical than conditional variance. The ratio of risk premium to variance is therefore weakly procyclical, unlike the Sharpe ratio. We document this fact in a broad global equity sample, using a range of methods to time portfolio formation after the onset of a recession. We also provide supporting evidence from option markets. We show that the ratio of risk premium to variance pins down the conditional beta in a regression of the stochastic discount factor on the market return, and its cyclicality is important for understanding stylized facts about the equity term structure. We present a stylized model that reconciles the procyclicality of the price per unit of variance risk with the term structure of Sharpe ratios to dividend claims.

*Jensen is at the Department of Finance, Bocconi University, Italy (christian.jensen@unibocconi.it), and Lazarus is at the Haas School of Business, University of California, Berkeley (lazarus@berkeley.edu). We are grateful to Niels Gormsen, Theis Ingerslev Jensen, Frank de Jong, Martin Lettau, Ian Martin, David Schreindorfer, and Tobias Sichert for helpful comments, as well as comments from seminar participants at the University of Copenhagen and Tilburg University. We thank Theis Ingerslev Jensen for sharing global returns data. The empirical results in Section 5 were originally reported in Lazarus (2022), which this paper now supersedes.

1. Introduction

The market's Sharpe ratio — the ratio of expected excess return to volatility — is well known to increase in bad times, when prices are low and volatility is high (Campbell and Cochrane (1999), Lustig and Verdelhan (2012)). In other words, the price per unit of *volatility* risk is at least somewhat countercyclical. The cyclicality in the market price per unit of *variance* risk, however, has received less attention in previous literature. Since volatility and variance map one-to-one, it is natural to conjecture that the cyclicalities in these unit prices of risks are similar. This conjecture turns out to be incorrect. We document that the price per unit of variance risk does not increase during economic downturns; in fact, across multiple specifications, it appears to be weakly procyclical rather than countercyclical. Dividing expected return by variance, rather than volatility, evidently matters, and this choice has meaningful implications for understanding risk pricing over the business cycle.

There are multiple possible notions of the price of risk. Why focus on the ratio of conditional risk premium to variance? We start with a set of general theoretical results motivating our analysis of this ratio, which we call γ_t. We show that this ratio pins down the conditional beta in a regression of the SDF onto the market return. Times with high γ_t therefore indicate that someone who is holding the market is exposed to greater SDF risk for every unit of market risk than normal. That is, a $\pm 1\%$ market return exposes the investor to greater SDF variation when the conditional beta is higher, and the ratio γ_t will be higher at these times as a result. We later use this insight to link the cyclicality of γ_t to the returns on equity index option portfolios at different horizons and the equity term structure.

We then turn to the data. We begin by constructing multiple measures of γ_t to examine its cyclical behavior. Since γ_t is not directly observable in the data, we start by showing that three fundamentally different approaches to computing the ratio have similar time-series patterns. Specifically, we compute (1) a realized measure from daily returns within each month, (2) an expected measure using the methods of Kelly and Pruitt (2013) that relies on the cross-section of valuation ratios, and (3) an option-implied measure computed from options written on the stock market. Time-series predictions for the ratio γ_t from these three approaches are positively correlated with each other, with pairwise correlations ranging between 0.27 and 0.42.

Using each of the three measures for the ratio γ_t, we next investigate their time-series properties in more detail. We start by focusing on the differences in the price per unit of risk that an investor can earn in normal times, which we here define as NBER non-recession periods, to what can be earned by an investor who enters the market during recessions. However, as discussed in detail in Lustig and Verdelhan (2012), recessions are typically determined ex post and are based on negative economic activity that often coincides with initial turbulence in financial markets. To circumvent the ex post bias related to the negative economic activity in the beginning of recessions, we ask what an investor can earn if she has perfect foresight in recessions, meant in the way that she can pinpoint in real-time when the market has reached its low. She buys the market at its low and holds the market for twelve months thereafter. As expected, the perfect foresight investor earns statistically significantly higher Sharpe ratios than the investor who holds the market in normal times, confirming previous results in Campbell and Cochrane (1999) and Lustig and Verdelhan (2012).

In contrast, we find that the price per unit of variance risk is not statistically different for the two investors in the full sample starting in 1926. In the shorter samples from 1964 and 1996, in which we have data to compute the expected and option-implied measures of γ_t, we find that the perfect foresight investor earns significantly lower price per unit of variance risk than the investor who holds the market in normal times. We confirm these results in a global sample covering 20 stock market indexes around the world using OECD recession indicators. We also show that an investor without perfect foresight in recessions who buys the market either one, six, or twelve months into each recession earns significantly lower γ_t than the investor who holds the market in normal times. This holds true despite the fact that γ_t has a strong, positive unconditional correlation with the market's Sharpe ratio. The key distinction is that the Sharpe ratio robustly increases in bad times, while the price per unit of variance risk does not (and decreases significantly in many specifications). As a result, we show that the Sharpe ratio is significantly more countercyclical than γ_t both statistically and economically.

To further strengthen our main empirical results, we investigate how γ_t varies with macroeconomic variables that are typically tied to the state of the economy. We find that γ_t is high when recession probabilities are low, when financial markets are experiencing looser-than-average conditions, and when the economy is in an expansionary state with increasing inflation. We also find that γ_t is positively related to future 1- and 2-year growth in consumption and growth in industrial production, suggesting that the price per unit of variance risk is high in good times. Lastly, we show that γ_t is contemporaneous negatively related to the dividend-price ratio, suggesting that the price per unit of variance risk is high at times when prices are high. We confirm these results in a global sample and with various measures for the price per unit of variance risk. The Sharpe ratio, by contrast, is much more weakly related to these variables.

Having thoroughly established that γ_t is indeed at least weakly procyclical in the data, we next turn to the implications for the equity term structure. We consider three different notions of the equity term structure: (i) the holding period risk premiums in a CAPM-type model, (ii) returns on option portfolios at different horizons, and (iii) the term structure of risk-adjusted returns to dividend claims. We start from a simple CAPM-type model to show the cyclicality in the one-period price per unit of variance risk is related to the term structure of holding-period returns. Adding a bit of structure on market variance, the price per unit of this variance risk, and their correlation, we show that there is a clear connection between how they correlate and how risk is priced at different horizons. Under reasonable parameter values that we take from the data, we show that a procyclical price per unit of variance risk implies an unconditional downward sloping term structure of holding period returns.

To further extend our term structure results, we next turn to the implications for the term structure of option portfolio returns at different horizons. We start by looking to option markets where we consider a binary bet with constant quantity of risk in that it pays off either 1 or -1 if certain return outcomes realize. We show that the one-period expected return on this bet depends only on the price per unit of variance risk. We then show that, if the price per unit of variance risk is negatively correlated to the SDF then a multiperiod bet will provide a hedge against shocks to the SDF, implying that multiperiod expected returns should be lower than one period returns. This implies a downward-sloping term structure of risk prices for option portfolios that fix the quantity of risk. We provide empirical evidence in favor of these insights. These results are, in effect, an out-of-sample test of the procyclicality of γ_t established earlier in the paper: we show that the particular option-implied term structure considered here is downward-sloping if and only if γ_t is procyclical, so the fact that we indeed find a downward-sloping term structure provides further support to the preceding results.

Finally, we present a stylized model that links the cyclicality in the price per unit of variance risk to the slope of the term structure of the Sharpe ratio of dividend claims. In our model, return volatility has three components: fundamental dividend volatility (which is constant), discount-rate volatility (which increases in the price of fundamental risk), and non-fundamental volatility (which also increases in the price of fundamental risk). An increase in the price of fundamental risk therefore increases "pure" market risk that is not fully connected to fundamentals, increasing market variance without passing through one-for-one to expected returns. While this is sufficient to generate a procyclical price per unit of variance risk that decreases in price of fundamental risk, this non-fundamental volatility effect is not strong enough to obtain a procyclical Sharpe ratio, i.e., Sharpe ratios are countercyclical in the model. Further, the greater exposure of long-maturity claims to discount-rate risk means that their volatility increases without changing their expected return, generating a downward-sloping Sharpe ratio of dividend claims.

As a result, this stylized model shows how our findings about the cyclicality of the price per unit of variance risk connect to facts about the dividend term structure. We obtain *both* (i) a lower beta of the SDF onto the market in bad times (our main stylized fact), and (ii) a downward-sloping Sharpe ratio of dividend claims by maturity, both through the same channel (non-fundamental return risk). Meanwhile, our parameterization maintains the usual countercyclical Sharpe ratio.

Related literature. Our empirical results on the cyclicality of the price per unit of variance risk relate most closely to the recent strand of literature that investigates the relationship between the market excess return or state prices and market variance.[1] Indeed, Moreira and Muir (2017) show that an investor can earn high Sharpe ratios by moving into the market when volatility decreases, suggesting that the price of risk is inversely related to volatility. Also, using option prices written on the market, Schreindorfer and Sichert (2023) show that the stochastic discount factor projected onto the market return is flatter at times when market volatility is high. Our results are consistent with the findings in this previous literature, but we extend the literature in several important dimensions. First, we do not only focus on market variance as a measure of the state of the economy. Instead, we investigate how the price per unit of market variance risk varies with the business cycle using a broad set of economic indicators that are often linked to the economic activity like recession indicators, valuation ratios, and growth in industrial production and consumption. We also extend the results to an international setting covering 20 stock markets around the world, showing that the procyclicality of the price per unit of variance risk is a world-wide phenomenon. Lastly, while the focus of the previous papers are on the ability of the leading asset pricing models to match their stylized facts, we focus on bridging the gap between the literature that investigates time variation in the price of risk with the equity term structure literature.

Our analysis of the term structure of option portfolio returns at different horizons is related to results in Bliss and Panigirtzoglou (2004) who conduct a similar analysis using option information but with stronger parametric testing assumptions. More generally, our option results on both the computation of the time varying γ_t and the unconditional option portfolio term structure relate to a longstanding literature that attempts to extract

[1] While our focus is primarily on the price per unit of variance risk, we reconfirm the countercyclical patterns in the Sharpe ratio discussed in Campbell and Cochrane (1999) and Lustig and Verdelhan (2012).

an implied market risk aversion from option prices, see e.g. Ait-Sahalia and Lo (2000), Jackwerth (2000), and Rosenberg and Engle (2002). Specifically, when computing γ_t from option prices, we extend the previous literature by combining the methodology of Bliss and Panigirtzoglou (2004) with a set of linear constraints from Jensen, Lando, and Pedersen (2019), which allows us to obtain a γ_t that is time varying.

The paper proceeds as follows. In Section 2, we motivate our analysis of the price per unit of variance risk. In Section 3, we discuss how we compute the price per unit of variance risk empirically. Section 4 presents our main empirical results on the cyclicality of the price per unit of variance risk. In Section 5, we document how the procyclical patterns of the price per unit of variance risk relate to the slope of the equity term structure. Section 6 concludes.

2. Theory

In this section, we theoretically motivate our subsequent analysis of the ratio

$$\gamma_t \equiv \frac{\mu_t}{\sigma_t^2} = \frac{\mathbb{E}_t[R_{m,t+1} - R_{f,t+1}]}{\text{Var}_t(R_{m,t+1})}, \tag{1}$$

where $R_{m,t+1}$ is the gross return on the market and $R_{f,t+1}$ is the gross risk-free rate. This ratio is equivalent to $\gamma_t = SR_t/\sigma_t$, where $SR_t \equiv \mu_t/\sigma_t$ is the conditional Sharpe ratio.

Assuming the absence of arbitrage, there exists a strictly positive one-period stochastic discount factor (SDF) M_{t+1} such that $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ for any gross return R_{t+1}. As is standard, one can use the definition of the conditional covariance to write this as

$$\mathbb{E}_t[R_{t+1} - R_{f,t+1}] = -R_{f,t+1}\text{Cov}_t(M_{t+1}, R_{t+1}). \tag{2}$$

It is common to rewrite (2) to obtain a single-beta representation for expected returns: $\mathbb{E}_t[R_{t+1} - R_{f,t+1}] = \beta_{R\to M,t}\lambda_{M,t}$, where $\beta_{R\to M,t} \equiv \frac{\text{Cov}_t(M_{t+1},R_{t+1})}{\text{Var}_t(M_{t+1})}$ is the slope in a regression of return R_{t+1} onto the SDF, and $\lambda_{M,t} \equiv -R_{f,t+1}\text{Var}_t(M_{t+1})$. According to this representation, assets differ only in their quantity of SDF risk $\beta_{R\to M,t}$, and there is a single SDF factor risk premium $\lambda_{M,t}$ that is often referred to as the *price of risk* (e.g., Cochrane, 2005; Campbell, 2018) for all assets.[2] While the common risk premium feature is appealing, this is of course not the only available representation for expected returns, nor the only

[2]Other notions of the price of risk also exist in other contexts. In continuous-time models with a single Brownian shock, for example, the *market price of risk* often refers to the Sharpe ratio for an asset exposed to that shock.

possible notion for the price of risk. We take a different — effectively converse — route.

In particular, we rewrite (2) as

$$\mathbb{E}_t[R_{t+1} - R_{f,t+1}] = -R_{f,t+1}\,\beta_{M\to R,t}\,\sigma_t^2, \tag{3}$$

where $\beta_{M\to R,t} \equiv \frac{\text{Cov}_t(M_{t+1},R_{t+1})}{\text{Var}_t(R_{t+1})}$ and $\sigma_t^2 \equiv \text{Var}_t(R_{t+1})$. In this representation, the intuitive labels separating price from quantity of risk are effectively reversed: one can think of σ_t^2 as the quantity of asset-specific risk, and the SDF's exposure to this asset, $\beta_{M\to R,t}$, is the price per unit of asset-specific risk.

Specializing (3) to the case of the market return and rearranging, we see that γ_t — the ratio of market risk premium to variance — pins down the loading in a regression of the SDF onto the market. We summarize this in the following result.[3]

Result 1. *The ratio of the risk premium μ_t to variance σ_t^2 satisfies*

$$\gamma_t = -R_{f,t+1}\beta_{M\to R,t},$$
$$\text{where} \quad \beta_{M\to R,t} \equiv \frac{\text{Cov}_t(M_{t+1}, R_{m,t+1})}{\sigma_t^2}.$$

The cyclical behavior of $\gamma_t = \mu_t/\sigma_t^2$ therefore speaks to the cyclical exposure of the SDF to the market. So while we are not the first to study the behavior of γ_t — among recent literature, we follow (and extend) Moreira and Muir (2017) most closely in doing so — the characterization in terms of $\beta_{M\to R,t}$ is novel, and we show below that it allows us to tie the behavior of γ_t to the equity term structure. The result tells us that increases in γ_t must mean that someone holding the market is exposed to greater SDF risk for every unit of market risk. That is, a $\pm 1\%$ market return exposes the investor to greater SDF variation when $\beta_{M\to R,t}$ is higher, and the ratio γ_t will be higher at these times as a result.

One can also relate $\beta_{M\to R,t}$ and γ_t to more classical notions of the price of risk, following the discussion around equation (3). If the market is mean–variance efficient (as, e.g., in the CAPM), then for an arbitrary asset,[4]

$$\mathbb{E}_t[R_{t+1} - R_{f,t+1}] = \gamma_t \text{Cov}_t(R_{t+1}, R_{m,t+1}).$$

Fama (1968) refers to γ_t as the "market price per unit of risk" as a result of the above

[3] Note that we slightly abuse notation in continuing to refer to this loading as $\beta_{M\to R,t}$, as above, even when specializing to the market return.

[4] To see this, note that mean–variance efficiency implies that $M_{t+1} = a - bR_{m,t+1}$ (Cochrane, 2005), so $\beta_{M\to R,t} = -b$. And from (2), $\mathbb{E}_t[R_{t+1} - R_{f,t+1}] = -R_{f,t+1}\text{Cov}_t(M_{t+1}, R_{t+1}) = bR_{f,t+1}\text{Cov}_t(R_{t+1}, R_{m,t+1}) = \gamma_t\text{Cov}_t(R_{t+1}, R_{m,t+1})$.

relationship. Alternatively, if returns are log-normal and the SDF is proportional to $R_{m,t+1}^{-\widetilde{\gamma}_t}$ — as would be the case, for example, with a representative agent with relative risk aversion $\widetilde{\gamma}_t$ facing i.i.d. consumption growth — then in fact $\gamma_t \approx \widetilde{\gamma}_t$.[5] Friend and Blume (1975), among others, refer to γ_t as the market price of risk as a result of a version of this observation; others who refer to relative risk aversion as synonymous with the price of risk are effectively doing the same.

We return to these characterizations of γ_t — both the characterization in terms of the SDF loading on the market in Result 1, and the characterization as effective market risk aversion — after presenting our empirical results, which we turn to now.

3. Inferring the conditional price per unit of variance risk

Before going into details about our empirical results, we start by describing how we compute the price of market variance risk, γ_t from equation (1). Since γ_t is not directly observable, we consider three fundamentally different approaches to infer γ_t.

Realized conditional price per unit of variance risk — $\gamma_t^{\textbf{realized}}$

Our first measure of the price per unit of variance risk is an ex-post measure that relies on realized within month daily returns. Let $\tilde{R}_s = R_s - R_s^f$ be the excess return on date s. We compute the realized conditional price per unit of variance risk in month t as

$$\gamma_t^{\text{realized}} = \frac{\sum_{s=1}^{N_t} \tilde{R}_s}{\frac{N_t}{N_t-1} \sum_{s=1}^{N_t} [\tilde{R}_s - (\sum_{s=1}^{N_t} \tilde{R}_s)]^2} \tag{4}$$

where N_t denotes the number of trading days in month t.[6]

Expected conditional price per unit of variance risk — $\gamma_t^{\textbf{expected}}$

Our second approach of computing the price per unit of variance risk relies on ex ante predicted values for the conditional market risk premium and its variance. We compute the conditional market risk premium using the methodology in Kelly and Pruitt (2013).

[5]Given $M_{t+1} \propto R_{m,t+1}^{-\widetilde{\gamma}_t}$, we have $\widetilde{\gamma}_t = -\text{Cov}_t(\log M_{t+1}, \log R_{m,t+1})/\sigma_t^2$, where σ_t^2 is now the variance of log returns. Thus $\widetilde{\gamma}_t \sigma_t^2 = -\text{Cov}_t(\log M_{t+1}, \log R_{m,t+1})$, which by the pricing equation (and under log-normality) is $\log \mathbb{E}_t[R_{m,t+1}] - \log R_{f,t+1} \approx \mathbb{E}_t[R_{m,t+1} - R_{f,t+1}]$. We note that $\widetilde{\gamma}_t$ in $M_{t+1} \propto R_{m,t+1}^{-\gamma_t}$ can also be viewed as a reduced-form variable corresponding to the as-if relative risk aversion over market returns as of time t (i.e., this representation is more general than the power-utility case). Another way to see this is to use the standard myopic portfolio choice rule, $w_m = \frac{\mathbb{E}_t[R_{m,t+1} - R_{f,t+1}]}{\widetilde{\gamma}_t \sigma_t^2}$, where w_m is the share of wealth invested in the market. Setting $w_m = 1$ in equilibrium, we again obtain $\widetilde{\gamma}_t = \gamma_t$.

[6]We thank Theis Ingerslev Jensen for sharing daily excess returns on international stock market indexes.

Under the two assumptions: (i) the expected log return and log growth rates are linear in a set of latent factors, and (ii) these factors evolve according to a first-order vector autoregression, Kelly and Pruitt (2013) show how to infer the conditional market risk premium from the cross-section of valuation ratios. The main reason why we choose this estimator as our ex ante predictor of the market risk premium is that the estimator does well in predicting market returns both in- and out-of-sample. Kelly and Pruitt (2013) find that they can predict the one-month market risk premium on the U.S. market portfolio with an R^2 of 2.38 in-sample and an R^2 of 0.93 out-of-sample. A minor benefit is that the method is build around valuation ratios and the predicted expected returns are therefore likely to fluctuate with the business cycle as we would expect.

Consistent with previous literature that consider time-variation in market variance (see e.g. Campbell et al. (2018)), we compute conditional expected market variance assuming that the variance follows a first-order autoregressive process. We compute the predicted variance via the relationship

$$\tilde{\text{Var}}_t(R_{m,t+1}) = \theta_0 + \theta_1 \text{Var}_{t-1}(R_{m,t}) \tag{5}$$

where

$$\text{Var}_{t-1}(R_{m,t}) = \frac{N_t}{N_t - 1} \sum_{s=1}^{N_t} [\tilde{R}_s - (\sum_{s=1}^{N_t} \tilde{R}_s)]^2 \tag{6}$$

is the realized variance in month t. We infer the values of the parameters θ_0 and θ_1 in equation (5) from a linear regression of realized variance on its one period lagged value.

Combining the predictions from Kelly and Pruitt (2013) about the conditional market risk-premium with the AR(1) variance prediction, we compute $\gamma_t^{\text{expected}}$ as in equation (1).

Option implied conditional price per unit of variance risk — γ_t^{option}

As the third an final approach for computing the price per unit of variance risk, we look to option markets.[7] The premise for this approach is that the projection of the stochastic discount factor onto the market return, $M_{t+1}|R_{m,t+1} = \delta_t R_{m,t+1}^{\gamma_t}$, prices the market and derivatives written on the market. This premise is common in previous option literature, see e.g. Bliss and Panigirtzoglou (2004). Under this premise, we can relate the state price density ($\pi_{m,t+1}(x)$) of market returns to the physical probability density ($p_{m,t+1}(x)$) and a risk adjustment in the following way:

$$\pi_{m,t+1}(x) = p_{m,t+1}(x)\delta_t x^{-\gamma_t} \tag{7}$$

[7] The empirical methodology in this section was first reported in Chapter 3 of Jensen (2018), which this paper now supersedes.

Using insights from Breeden and Litzenberger (1978), we can use option prices to back out risk-neutral densities, say $f_t^*(R_{m,t+1}) = \pi_{m,t+1}(x)R_t^f$. These densities reflect the time t real-time risk-adjusted probabilities over the potential future market outcomes.

Now, from Equation (7), we can write the stock market's physical probability distribution function, say $F_{m,t}(x)$, as

$$F_{m,t+1}(x) = \int_{-\infty}^{x} p_{m,t+1}(y)dy = \int_{-\infty}^{x} \frac{\pi_{m,t+1}(y)y^{-\gamma_t}}{\delta_t}dy \tag{8}$$

If we knew the true values of the parameters δ_t and γ_t then we could directly infer the stock market probability distribution from the observable state price density. However, the true values of the parameters are not directly observable and we therefore have to come up with a way to infer them. To achieve this task and infer the true values of δ_t and γ_t, we follow Bliss and Panigirtzoglou (2004) and use the so-called Berkowitz test, cf. Berkowitz (2001). The idea behind the Berkowitz test is that, for the true values of δ_t and γ_t, the distribution of $u_{t+1} = F_{m,t}(R_{m,t+1})$ is uniform and the distribution $y_{t+1} = \Phi^{-1}(u_{t+1})$ is standard normal. Therefore, to conduct the Berkowitz test, we estimate the coefficients in the regression model:

$$y_{t+1} = a + \beta y_t + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, \sigma^2) \tag{9}$$

and perform a likelihood ratio test of the joint hypothesis that $a = \beta = 0$ and $\sigma^2 = 1$.[8] It is worth noticing that, even though there might be momentum effects in returns, then we will still want $b = 0$ because the true distribution should take these momentum effects into account. The Berkowitz likelihood ratio test for non-overlapping returns is:

$$LR = -2(LL(0,0,1) - LL(a, \beta, \sigma^2)) \sim \chi_3^2 \tag{10}$$

where $LL(a, \beta, \sigma^2))$ is the log likelihood of Equation (9). The likelihood ratio test statistic, LR, is chi-square distributed with three degrees of freedom.

To find the values of δ_t and γ_t, we minimize the Berkowitz test statistic in Equation (10) under the constraint that, for all dates t, the equation $\int_{-\infty}^{\infty} \frac{\pi_t(y)y^{-\gamma_{m,t}}}{\delta_{m,t}}dy = 1$ must hold. This constraint ensures that the resulting physical return distributions integrate to one at

[8] We use non-overlapping monthly horizon distributions and returns. The hypothesis that $b = 0$ is therefore natural. For the case with overlapping returns see e.g. Bliss and Panigirtzoglou (2004) for a thorough discussion of the test.

all points in time. Written in mathematical terms, the optimization problem is:

$$\min_{\delta_t} -2\left(LL(0,0,1) - LL(a,\beta,\sigma^2)\right) \tag{11}$$

$$s.t. \quad \gamma_t \text{ solves } \int_{-\infty}^{\infty} \frac{\pi_t(y)y^{-\gamma_t}}{\delta_t} dy = 1, \quad \text{for all } t \tag{12}$$

For a given value of δ_t, the constraints provide enough equations to solve for the time-varying γ_t. Specifically, for a given level of δ_t, at any point in time, we only have to solve for γ_t. If γ_t was linear in the constraint, then solving for the parameter would be straightforward. However, γ_t enters non-linearly in the constraint and we need to address this non-linearity. The generalized recovery methodology of Jensen, Lando, and Pedersen (2019) provides us with the argument we need. If we assume that there is a solution to the constraint for one given γ_t, then that solution is almost surely unique. Practically, this means that there will be at most one solution to the constraint equation.[9]

To optimize over the parameter, δ_t, we need to make an assumption on its functional form. We allow δ_t to be time varying through the time variation in the gross risk-free rate. Specifically, we assume that $\delta_t = \frac{1}{R_t^f} + c$ where R_t^f is the time t gross risk-free rate and c is a time-invariant parameter. This functional form of $\delta_{m,t}$ is conveniently simple while nesting the risk-neutral distribution as a solution if $c = 0$. To minimize Equation (11), we search over a grid of values for c and pick the c which provides the lowest Berkowitz test statistic. Importantly, for each value of c, the constraints ensure that we can infer a time varying level of γ_t. For different values of c, the γ_t time series will differ and consequently also the physical distributions, which gives us the variation in the Berkowitz test statistics that we need for our optimization.

We set γ_t^{option} to be the optimized γ_t from the optimization problem in 11. That is, γ_t^{option} takes the values that best reconciles the ex ante observable option prices with the ex post realized returns.

The Sharpe ratios. We compute the Sharpe ratios by multiplying the estimated price per unit of variance risks by the conditional volatility. For the realized Sharpe ratio, we multiply $\gamma_t^{realized}$ with the within month standard deviation of returns. For the expected Sharpe ratio, we multiply $\gamma_t^{expected}$ with the expected volatility from the AR(1) model. Lastly, for the option Sharpe ratio, we multiply γ_t^{option} with the option implied volatility

[9] Our methodology is closely related to the methodology used in Bliss and Panigirtzoglou (2004). In short, the difference in the two methodologies is that, they optimize over a constant level of stock market γ whereas we optimize over the stock market time preference parameter δ_t and allow γ_t to be time-varying through the constraints in Equation (11).

computed from equation (7) using γ_t^{option} as the exponent.

Pairwise correlations. Figure 1 shows the estimated Sharpe ratio in panel (a) and prices per unit of variance risk in panel (b) for the US stock market. As reported in Table 1, the estimates for the prices of risks are highly correlated with correlations ranging from 0.27 to 0.42. In the table, we also report correlations for γ_t^{PCA} and SR_t^{PCA}. These are the first principal components for the price per unit of variance risk measures and the Sharpe ratio measures respectively. The first principal components capture a large part of the variation in the three measures with individual correlations to the measures ranging from 0.41 to 0.80. These principal components suggests that there is a strong similarity in the variation in the measures.

4. Cyclicality of risk prices

In this section, we investigate how γ_t varies over the business cycle. We start by investigating how risk prices change during recessions. Thereafter, we investigate if risk prices are different in and out of recession periods. Lastly, we investigate how the risk prices vary with macroeconomic variables that are typically tied to economic activity. In the following, we devote a subsection to each analysis.

4.1. Do risk prices increase from the onset of a recession to the end of the recession?

Figure 2 sets the stage for this analysis. The figure shows risk prices implied from options written on the US stock market index. The dashed black line is the implied Sharpe ratio of market returns and in the dashed blue line we plot the implied price per unit of variance risk. In the figure, we standardize both measures to make them comparable. A value of one should therefore be interpreted as a value that is one standard deviation above the average level for the measure. Both risk prices are clearly time varying and move together in a linear sense such that their correlation is 0.54 in our sample. Focusing on the shaded areas, which represent NBER recession months, the arrows show the increase or decrease in the risk prices from the onset of the recession to the end of the recession. The black arrows are the Sharpe ratio and the blue arrows are the price per unit of variance risk. In all three recessions in this sample, we find that the Sharpe ratio of market returns increased relatively more than the price per unit of variance risk. Actually, in all recessions the Sharpe ratio increased while the price per unit of variance risk decreased in the more recent covid period. Focusing on the financial crisis, the Sharpe ratio increased by almost two standard deviations, which corresponds to an annualized increase in the Sharpe ratio of 0.34. Looking at Figure 2, we conjecture that the Sharpe ratio increases both statistically and economically in recessions while the price per unit of variance risk does not.

To strengthen our conjecture, we next test if the Sharpe ratio indeed increases in recessions. The first and second rows of Table 2 report the results of this analysis for the Sharpe ratio. Using either of our three measures, we compute the difference in the Sharpe ratio from the onset of each recession to the end of the recession. The table reports t-tests for the hypothesis that the difference is zero. Using either of the three measures, we find that the difference is on average positive, ranging from 0.58 to 1.28, which corresponds to an annualized average increase of 0.10 to 0.22 in the Sharpe ratio during recessions. For the realized and the expected measures, this difference is statistically significant. These are also the measures for which we have the most observations. For the option measure, the t-stat is 2.80 and the 10% critical value with only three observations is 2.92. In the last column, we report the results using the first principal component of the three measures. Here, we find that the difference is statistically significant even though we have only a few observations.

The third and fourth columns of Table 2 report the results for the price per unit of variance risk. For this risk price, we find that the difference is mildly positive in two specifications and it is even negative for the option measure. Only for the realized measure is the risk price statistically positive. These results suggest to us that the price per unit of variance risk does not go up from the onset of a recession to the end of the recession.

Lastly, columns five and six report tests for the difference in the standardized changes to the Sharpe ratio and the price per unit of variance risk. Here we find that the difference in the changes int he risk prices during recessions is statistically significant in three of our four speifications and the point estimate is positive in all of them. These results speak to the cyclicality of the risk prices in that the Sharpe ratio is countercyclical in the sense that it increases during recessions from its pre-recession values while the price per unit of variance risk is acyclical in that there are only mild differences from the onset of a recession to the end of a recession in this risk price.

4.2. Are risk prices highest in recessions?

Recessions are typically determined ex post and are defined by a significant decline in economic activity. During such periods, we often observe large declines in industrial production, employment, and gross domestic product. In many recessions, we also observe initial large declines in the stock market. This is particularly the case if the recession was driven or initiated by financial activity as in the Great Financial Crisis of 2007-2009.

Directly comparing the price of risk for an investor who invests in normal times, non-recession periods, to the price of risk for an investor who invests during the ”full” recessions therefore suffers from an ex post bias in the sense that recessions are ex post determined by the decline in economic activity. However, following the ideas of Lustig and Verdelhan (2012), we can ask how the price of risk in normal times relate to the price of risk in recessions where the investor enters the market at specific points during the recessions. To this end, we start from the following narrative. Suppose an investor wants to enter the market in a recession. If the investor could choose that point in time ex post of the recession, when would the investor enter the market? The natural answer is, the investor would choose to enter the market when the market has reached its low point in the recession. This investor is our starting point. We say that this investor has ”perfect foresight” in that she can pinpoint when the market is at its low in recessions. This investor enters the market at its low and holds the market for twelve consecutive months thereafter.

Table 3 reports the results for the Sharpe ratio for the investor with perfect foresight during NBER recessions and for the investor who invests in normal times, which for this table is NBER non-recession months. We compute the statistic for three different samples: (i) the full sample going back to 1926, (ii) a post 1964 sample, and (iii) a post 1996 sample. We compute the unconditional Sharpe ratio for the investor with perfect foresight in recessions by bundling the monthly excess returns over all the recessions, the twelve months after the low in each recession, and computing the unconditional expected excess return and unconditional volatility using these monthly returns. The first part of the table reports the results for the unconditional Sharpe ratio and our results support evidence from Lustig and Verdelhan (2012) that, at least in some specifications, the Sharpe ratio is higher in recessions than in normal times. We find that this is true in all three samples and the difference is statistically significant in the two longest samples. Standard errors are bootstrapped using 10.000 samples.

The second part of Table 3 reports the average realized conditional monthly Sharpe ratio for the investors. We compute the conditional Sharpe ratios using within month daily excess returns to infer the conditional expected excess return and the conditional volatility. The average conditional Sharpe ratio is higher for the investor with perfect foresight and the difference is statistically significant in the two longest sample. The third and fourth part of the table report results for the conditional expected Sharpe ratio and the conditional option implied Sharpe ratio. The expected Sharpe ratio has on average lower Sharpe ratios for the perfect foresight investor but the option implied measure is positive and statistically significant at the 10% level. Overall, the results presented in Table 3 support the notion that the Sharpe ratio is countercyclical.

Next, we turn to the price per unit of variance risk. These results are reported in Table 4. The first part of the table reports the result for the unconditional price per unit of variance risk. Similarly to how we compute the unconditional Sharpe ratio, we bundle returns for the perfect foresight investor who invests after the market low in each recession and holds the market for twelve months thereafter. We then compute the unconditional price per unit of variance risk as the ratio of the average excess returns over the variance of these returns. In the longer samples, we find that the difference is positive and statistically significant, meaning that in these samples we find that the price per unit of variance risk is unconditionally higher than in normal times.

Turning to the conditional prices per unit of variance risk, we find similar results for the long sample and for the realized measure. Here the coefficient is positive and statistically significant at the 10% level. In all other conditional specifications, for realized, expected, and option implied measures and various sample lengths, we find that the point coefficient is negative and it is statistically significant in three of the five specifications. These results suggest that conditional price per unit of variance is lower in recessions than in normal times. Overall, our results on the price per unit of variance risk for the perfect foresight investor suggests that it is likely procyclical or in some specifications only weakly countercyclical or acyclical.

However, most investors do not have perfect foresight and might not move into the market at its exact low in recessions. A more realistic and implementable strategy is to enter the market a certain number of months into each recession, similarly to the real time method of Lustig and Verdelhan (2012). Since the more realistically implementable investment strategies do not necessarily enter the market at its low, it is likely that they will underperform the investor with perfect foresight and therefore have lower risk prices. Next, we investigate if this statement holds true by comparing how different implementable strategies perform relative to the normal times investor.

Table 5 reports the differences in the Sharpe ratio for the investor without perfect foresight that enters the market six months into each recession and holds the market for twelve months thereafter to the Sharpe ratio for the normal times investor. We find that the Sharpe ratio is highest for the normal times investor in six of our nine specifications and it is statistically significantly negative in two of these tests. These results suggest that for the more realistic trading strategy where the investor enters the market six months into each recession, the Sharpe ratio that the investor earns is largely indistinguishable for what is earned in normal times, if anything it is likely to be lower than what the normal times investor earns. These results remain if we consider other timings like one or twelve months into each recession instead of six months.

When looking at Table 6, that reports the results for the price per unit of variance risk, we find that this risk price is negative in eight of the nine specifications and it is statistically significant in five of these. These are all the specifications fo rthe conditional price per unit of variance risk in the modern sample starting in 1964. Our results for the price per unit of variance risk show that the investor who invests in normal times earn significantly higher risk prices than the investor who enters six months into each recession. As for the Sharpe ratio, these results remain for other timing windows for the investor who enters in recessions.

The previous tables report the results for our US sample. Next, we broaden to an international setting where we have recession data from the OECD database. Merging this data with our return data, we end up with a sample of 20 stock market indexes around the world, including the OECD sample for the US.

Table 7 reports the pooled sample differences in the monthly prices of risk during normal times, which for this table is OECD non-recession months, and the recession period trading strategies from Tables 3 to 6, using OECD recession indicators. The first row reports the results for the investor who has perfect foresight in recessions. The results in the first columns show that if we can perfectly pinpoint the low in recessions around the world, then we can earn higher Sharpe ratios in recessions than out of recessions, consistent with previous results in Campbell and Cochrane (1999) and Lustig and Verdelhan (2012). The second and third column confirms the results of Table 6, that the investor with perfect foresight in recessions earn a price per unit of variance risk that is similar to that of the normal times investor. The remaining rows confirm that realistically implementable trading strategies have lower risk prices in recessions.

To sum up, Tables 3 through 7 provide the following insights. Investors who can perfectly time the market in recessions, that is, invest at its low, can earn Sharpe ratios that are higher than those earned in normal times. However, the price per unit of variance risk earned by the investor with perfect foresight is at best equal to what is earned in normal times. This last finding, which is novel to this paper, holds for both conditional and unconditional measures of price per unit of variance risk, is robust to different definitions of recessions (NBER and OECD indicators), and holds for a pooled international sample. Importantly, we also find that realistic trading strategies that enter the market in recession periods earn lower price per unit of variance risk than what is earned in normal times while the results for the Sharpe ratio are largely indistinguishable for realistic trading strategies in recessions and the normal times invsetor.

Next, we investigate how the conditional risk prices vary with macroeconomic variables that capture the state of the economy.

4.3. Do risk prices move with the business cycle?

In this section, we further study the cyclical fluctuations risk prices by linking their fluctuations to variables that capture the state of the financial sector and the overall macro economy. As in the previous section, we first focus on our US sample and thereafter extend the results to a broader international setting.

We start by considering the relation between risk prices and Chicago Fed financial and macroeconomic indicators. The NFCI is the national financial condition indicator, which is comprised of several subcategories built to capture risk, credit conditions, and financial and non-financial leverage. High values of the variables are historically associated with tighter-than-average conditions in financial markets, i.e., bad times. The first five rows of Table 8 report the results of regressions on the form:

$$\text{Price of risk}_t = \alpha + \beta \times \text{Financial Risk Indicator}_t + \epsilon_t \tag{13}$$

The results for the Sharpe ratio, which are shown in the first three columns of Table 8, we find that the realized and expected measures are generally negatively related to the financial indicators while the option measure is positively related to most indicators but the positive coefficients are not statistically significant. The last three columns of the table reports the results for the price per unit of variance risk. Here, we find that all measures are generally negatively related to the financial risk indicators and ten of the fifteen specifications are statistically significant. These results suggest that the price per unit of variance risk moves strongly procyclically with these risk indicators while the Sharpe ratio moves only mildly with the risk indicators.

To extend the recession results from the previous tables, in row six of Table 8, we also report results of the relationship between risk prices and the recessions probability of Chauvet and Piger (2008). We find that, when the probability of a recession is high then the price per unit of variance risk is low, lending further evidence of its procyclicality. Results are generally weaker for the Sharpe ratio.

In row seven, we extend the results to the Chicago Fed National Activity Index, CFNAI, which is an index build to capture overall economic activity and inflationary pressure. A high value of the CFNAI is generally associated with good economic conditions with high consumption growth, low unemployment, and high industrial production. We find that price per unit of variance risk is positively related to CFNAI and the effect is statistically significant in tho of the three measures of the price per unit of variance risk. For the Sharpe ratio, we find no clear relationship between any measure and the CFNAI.

Next, we turn to consumption growth. Due to the fact that in the US we have monthly data and in our international sample we have quarterly observations, we divide our analysis into two parts, a US part and an international part. We obtain data on US consumption from the St. Louis Fed database on monthly personal consumption expenditures of non-durable and service goods. We deflate consumption with the CPI. The last column of Table 8 reports the results when regressing the risk prices onto the future eight month consumption growth. Both for the Sharpe ratio and the price per unit of variance risk, we find only mild positive relationships between the measures and consumption growth for the realized and expected measures while the option measures have negative but insignificant coefficients..

In Table 9, we move to an international setting where in the first part, we report the results of a pooled panel regression of the risk prices onto the subsequent eight quarter consumption growth ($s = 4$ or $s = 8$):

$$\text{Risk price}^i = \alpha_i + \beta \times \text{consumption growth}^i_{q+1,q+s} + \epsilon^i_{q+1,q+s} \tag{14}$$

where i represent the different countries and Risk price^i is the average of the monthly measures within quarter q. We cluster standard errors by country and quarter and include country fixed effects. We use Final Consumption Expenditure, Real, Unadjusted, Domestic Currency from the IMF database as our proxy for aggregate consumption. The first row reports the results when we pool the raw data and the second row reports results where we standardize both risk prices and consumption growth within each country before pooling the data. In all panel regression specification, we find that the price per unit of variance risk is positively related to future consumption growth and the slope coefficients are statistically significant for the realized measure. For the Sharpe ratio, we find similar results.

We next study how the risk prices vary with valuation ratios, which are standard measures of the state of the economy in previous asset pricing literature (see e.g. Campbell and Cochrane (1999) and Gormsen and Jensen (2024)). We measure valuation ratios through country-level dividend-price ratios and book-to-market ratios. In the second aprt of Table 9, we report the results of panel regressions on the form:

$$\text{Risk price}^i = \alpha_i + \beta \times \text{valuation ratio}^i_t + \epsilon^i_t \tag{15}$$

where we regress the risk prices onto the contemporaneous valuation ratio. We include country fixed effects and cluster standard errors by country and time. The first row for the dividend-price section reports the panel regression results when we pool the raw data for all countries. The second row reports the results where we standardize (mean zero and variance one) both the risk prices and the valuation ratio within each country before pooling. We find that the price per unit of variance risk is negatively related to the contemporaneous level of the dividend-price ratio, suggesting that investors can earn higher price per unit of variance risk in good times when market prices are high. For the Sharpe ratio, we find similar results, but they are mildly weaker for the expected measures.

Finally, we also consider how the risk prices vary with the growth in industrial production. The third part of Table 9 reports panel regressions on the form:

$$\text{Risk price}^i = \alpha_i + \beta \times \text{Industrial production}^i_{t+1,t+s} + \epsilon^i_t \tag{16}$$

where $s = 8$ are months. The first row of the industrial production part of the table reports results using the raw data and the second row reports the results when we standardize the data input within country before pooling the data. Our data on industrial production is from the OECD database. We find that the price per unit of variance risk is positively related to future growth in industrial production in all our regression specifications and the slopes are all statistically significant. The coefficients are also positive for the Sharpe ratio but less statistically significant for the expected measure. In these panel regressions, we again add country fixed effects and cluster standard errors by country and time. These results show that the risk prices are high in good times when economic activity as measures by the growth in industrial production is high.

5. Implications for the equity term structure

We now proceed to show that our results on the cyclicality of the price per unit of variance risk, γ_t, have implications for the behavior of the equity term structure. More specifically, we consider three separate notions of the term structure of equity claims. We start by investigating the unconditional holding period risk premiums in a CAPM-type model where market variance and the price per unit of variance risk correlate. Next, we show that our results have direct implications for the returns on equity index option portfolios at different horizons, which we provide preliminary evidence for in the data. Lastly, we show how our results relate to — and can be reconciled with — facts about the term premium on dividend claims at different horizons.

5.1. Holding period risk premiums in a CAPM-type model

This subsection serves the purpose of showing that even in a simple CAPM-type model it is natural to connect the cyclicality in γ_t, the price per unit of variance risk, to the unconditional equity term structure. In this subsection, we specifically look at term structure of *holding period* risk premiums, which is not the same as the term structure of one-period returns to dividend claims that are often discussed in the literature (e.g. van Binsbergen, Brandt, and Koijen (2012)). We turn to the returns on dividend claims later in the section.

We start from the common, Cochrane (2005), notation of the CAPM stochastic discount factor:

$$M_{t,t+1} = A_{t,t+1} - B_{t,t+1} R_{t,t+1} \tag{17}$$

where $R_{t,t+1}$ is the gross return on the market, $A_{t,t+1} = 1/R^f_{t,t+1} + B_{t,t+1}\mathbb{E}_t[R_{t,t+1}]$, $B_{t,t+1} = \frac{\mathbb{E}_t[R_{t,t+1} - R^f_{t,t+1}]}{\sigma^2_{t,t+1}} / R^f_{t,t+1} = \gamma_{t,t+1}/R^f_{t,t+1}$, and $R^f_{t,t+1}$ is the one period gross risk-free rate. Note that the "slope" parameter, $B_{t,t+1}$, is directly related to the price per unit of variance risk (not the Sharpe ratio).

Given the stochastic discount factor in (17), we can write the conditional expected excess market return in the usual way:

$$\mathbb{E}_t[R_{t,t+1} - R^f_{t,t+1}] = -R^f_{t,t+1} cov_t(R_{t,t+1}, M_{t,t+1}) \tag{18}$$

$$= \gamma_{t,t+1}\sigma^2_{t,t+1} \tag{19}$$

In a similar fashion, we assume that there is a representation of the two period market risk premium as:

$$\mathbb{E}_t[R_{t,t+2} - R^f_{t,t+2}] = \gamma_{t,t+2}\sigma^2_{t,t+2} \tag{20}$$

Here, $\gamma_{t,t+2}$ is the time t price per unit of variance risk over the period until time $t+2$.

We now write the unconditional market risk premium at different horizons as:

$$\mathbb{E}[R_{t,t+1} - R^f_{t,t+1}] = \mathbb{E}[\gamma_{t,t+1}]\mathbb{E}[\sigma^2_{t,t+1}] + \text{Cov}[\gamma_{t,t+1}, \sigma^2_{t,t+1}] \tag{21}$$

$$\mathbb{E}[R_{t,t+2} - R^f_{t,t+2}] = \mathbb{E}[\gamma_{t,t+2}]\mathbb{E}[\sigma^2_{t,t+2}] + \text{Cov}[\gamma_{t,t+2}, \sigma^2_{t,t+2}] \tag{22}$$

At first glance, it seems natural that the covariance between the price per unit of variance risk and the market variance is important for the unconditional term structure of equity returns.

Now, to make progress, we impose structure on σ_t^2 and $\gamma_{t,t+1}$:

$$\sigma_{t,t+1}^2 = a + \rho\sigma_{t-1,t}^2 + \epsilon_{t,t+1} \tag{23}$$
$$\gamma_{t,t+1} = \gamma + \theta\gamma_{t-1,t} + b\sigma_{t,t+1}^2 + \nabla_{t,t+1} \tag{24}$$

Variance is an AR(1) and the price of risk is an AR(1) augmented with a component that makes variance and the price of risk correlated ($b\sigma_{t,t+1}^2$). If $b < 0$ then they are negatively correlated since all shocks ($\nabla_{t,t+1}$ and $\epsilon_{t,t+1}$) are independent (with each other and over time).

Using equation (24) iteratively and assuming that both the market variance and the price per unit of variance risk are stationary, we arrive at an expression for the one period unconditional risk premium as:[10]

$$\mathbb{E}[R_{t,t+1} - R_{t,t+1}^f] = \mathbb{E}[\gamma_{t,t+1}]\mathbb{E}[\sigma_{t,t+1}^2] + \frac{b}{1-\theta\rho}\text{Var}(\sigma_{t,t+1}^2) \tag{25}$$

Turning to the two period unconditional risk premium, we assume that returns are uncorrelated over time such that the two period variance is the sum of one period variances:

$$\sigma_{t,t+2}^2 = \sigma_{t,t+1}^2 + \sigma_{t+1,t+2}^2 \tag{26}$$

We next set $\gamma_{t,t+2}$ to be the predicted value one period ahead:

$$\gamma_{t,t+2} \equiv \mathbb{E}_t[\gamma_{t+1,t+2}] = \gamma + \theta\gamma_{t,t+1} + b\mathbb{E}_t[\sigma_{t+1,t+2}^2] \tag{27}$$
$$= \gamma + \theta\gamma_{t,t+1} + b[a + \rho\sigma_{t,t+1}^2] \tag{28}$$

This choice implies that the unconditional term structure of the price per unit of variance risk is flat, similar to that of the unconditional variance.

Inserting (26) and (28) into the covariance in equation (22), we arrive at an expression for the "annualized" (divided by 2) two period risk premium as:

$$\mathbb{E}[R_{t,t+2} - R_{t,t+2}^f]/2 = \mathbb{E}[\gamma_{t,t+1}]E[\sigma_{t,t+1}^2] + b(1+\rho)\left[\rho + \frac{\theta}{1-\theta\rho}\right]\text{Var}(\sigma_{t,t+1}^2)/2 \tag{29}$$

since $\gamma_{t,t+2} \equiv \mathbb{E}_t[\gamma_{t+1,t+2}]$ and the unconditional expectation is $\mathbb{E}[\gamma_{t+1,t+2}] = \mathbb{E}[\gamma_{t,t+1}]$. We furthermore have that $\mathbb{E}[\sigma_{t,t+2}^2] = \mathbb{E}[\sigma_{t,t+1}^2 + \sigma_{t+1,t+2}^2] = 2\mathbb{E}[\sigma_{t,t+1}^2]$. With this expression,

[10] Note that the unconditional covariance becomes an infinite sum: $\text{Cov}[\gamma_{t,t+1}, \sigma_{t,t+1}^2] = b\text{Var}(\sigma_{t,t+1}^2)\sum_{n=0}^{\infty}(\theta\rho)^n = \frac{b}{1-\theta\rho}\text{Var}(\sigma_{t,t+1}^2)$ when $|\theta| < 1$ $|\rho| < 1$.

we can write the difference in the "annualized" holding period risk premium as:

$$\mathbb{E}[R_{t,t+2} - R^f_{t,t+2}]/2 - \mathbb{E}[R_{t,t+1} - R^f_{t,t+1}] = b(1+\rho)\left[\rho + \frac{\theta}{1-\theta\rho}\right]\text{Var}(\sigma^2_{t,t+1})/2 - \frac{b}{1-\theta\rho}\text{Var}(\sigma^2_{t,t+1}) \tag{30}$$

This equality leads us to Result 2:

Result 2 (Cyclicality in $\gamma_{t,t+1}$ and the slope of the term structure)**.** *The equity term structure is:*

(i) downward sloping if:

$$b\left[(1+\rho)[\rho + \theta(1-\rho^2)] - 2\right] < 0 \tag{31}$$

(ii) upward sloping if:

$$b\left[(1+\rho)[\rho + \theta(1-\rho^2)] - 2\right] > 0 \tag{32}$$

(iii) flat if:

$$b\left[(1+\rho)[\rho + \theta(1-\rho^2)] - 2\right] = 0 \tag{33}$$

The proof is in the body of the text above.

Result 2 highlights that the time varying relationship between market variance and the price per unit of this variance risk is important for the term structure of holding period returns. In the data, we find that $b < 0$, that is, $cov(\gamma_{t,t+1}, \sigma^2_{t,t+1}) < 0$. So for the term structure to be downward sloping, for example, we need $(1+\rho)[\rho+\theta(1-\rho^2)] > 2$. This inequality holds if θ and ρ are "large". For example, if $\theta = 0.90$, and $\rho = 0.75$. Loosely speaking, the inequality in equation (31) holds when $\rho\theta > 0.7$. Looking at the persistence in the data, we find that $\theta = 0.91$ (using the expected price per unit of variance risk measure described in Section 3) and $\rho = 0.61$. The persistence in the variance largely depends on the sample, when including large spikes like in the financial crisis of 2008-2009 and the Covid-19 period, the persistence tends to be lower and the R^2 of a simple regression of variance onto its lagged value tends to go down relative to what we find when excluding these extremes. A more realistic specification of the variance process that features jumps or non-linear terms, like a leverage effect, should be able to capture these extreme periods. However, this added complexity is outside the scope of our analysis.

Overall, Result 2 is important because it highlight that, even in a simple CAPM-type model with little structure on market variance the price per unit of this variance risk, we can clearly connect the cyclicality in the one period risk price to how risk is priced at different horizons.

5.2. Option-implied risk prices by horizon

Next, we turn to option markets and show that the term structure of returns to particularly interesting option portfolios is directly related to the cyclciality in the price per unit of variance risk. We start by going through our theoretical setting and we present empirical evidence thereafter.

Theory

To simplify exposition, we consider an economy in which returns and the SDF are conditionally jointly log-normal, with

$$r_{m,t+1} \equiv \log R_{m,t+1} = \mu_{R,t} + \sigma_{\varepsilon,t}\varepsilon_{t+1} - \frac{1}{2}\sigma_{\varepsilon,t}^2, \tag{34}$$

$$m_{t+1} \equiv \log M_{t+1} = -r_{f,t} - \gamma_t r_{m,t+1} + \sigma_{\eta,t}\eta_{t+1} - \left[\frac{1}{2}(\sigma_{\eta,t}^2 + \gamma_t(1+\gamma_t)\sigma_{\varepsilon,t}^2) - \gamma_t\mu_{R,t}\right], \tag{35}$$

where ε_{t+1} and η_{t+1} are standard normal and independent (both over time and with respect to each other). The last terms in both lines are Jensen's inequality corrections to ensure that $\log \mathbb{E}_t[R_{m,t+1}] = \mu_{R,t}$ and $\log \mathbb{E}_t[M_{t+1}] = -r_{f,t}$. Writing the unexpected part of m_{t+1} as a linear combination of a projection onto the market and an orthogonal term is without loss of generality in this setting. Define $\sigma_t^2 = \text{Var}_t(r_{m,t+1})$. Since $\mu_{R,t} - r_{f,t} = -\text{Cov}_t(m_{t+1}, r_{m,t+1})$, it follows that $(\mu_{R,t} - r_{f,t})/\sigma_t^2 = \gamma_t$. This motivates our use of γ_t to refer to the loading of the log SDF onto the market.

We now consider options. No arbitrage implies the existence of a risk-neutral density $f_t^*(R_{m,t+1})$ such that

$$f_t^*(R_{m,t+1}) = R_{f,t}\mathbb{E}_t[M_{t+1}|R_{m,t+1}]f_t(R_{m,t+1}), \tag{36}$$

where $f_t(R_{m,t+1})$ is the objective physical density. As observed by Schreindorfer and Sichert (2023), $f_t^*(R_{m,t+1})/R_{f,t}$ can be thought of as the price of the Arrow-Debreu security that pays 1 if the return $R_{m,t+1}$ is realized and 0 otherwise, while $f_t(R_{m,t+1})$ can be thought of as its expected payoff.[11] This implies that the log expected return of the Arrow-Debreu security is $-\log \mathbb{E}_t[M_{t+1}|R_{m,t+1}]$. Using the characterization in (34)–(35) and the assumption of log-normality, this implies that the log expected excess return is

[11]This is loose only insofar as these are continuous densities. To formalize this fully, one can either consider a discretized version of the state space (as we will do in the empirical analysis) or define the Arrow-Debreu (AD) payoff to be a Dirac delta function.

equal to

$$-\log \mathbb{E}_t[M_{t+1}|R_{m,t+1}] - r_{f,t} = \gamma_t r_{m,t+1} - \gamma_t \left(\mu_{R,t} - \frac{1}{2}\sigma^2_{\varepsilon,t} \right). \tag{37}$$

Now consider a strategy that goes short one unit of the AD security for return state $R_{m,t+1} = \omega_1$, and long one unit of the AD security for state $R_{m,t+1} = \omega_2$, where $\omega_2 > \omega_1$. This strategy can be thought of as a conditional binary bet: it involves a payoff of 1 if $R_{m,t+1} = \omega_2$, a payoff of -1 if $R_{m,t+1} = \omega_1$, and 0 otherwise. Denote the return on this strategy by $R_{\omega,t+1}$. Using (37), the log expected return on the strategy is

$$\begin{aligned} \log \mathbb{E}_t[R_{\omega,t+1}] &= \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_1] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_2] \\ &= \gamma_t(\omega_2 - \omega_1). \end{aligned} \tag{38}$$

Intuitively, γ_t is the market price per unit of variance risk. This strategy fixes the quantity of risk: it pays off either 1 or -1, and importantly, it holds fixed the return outcomes ω_1 and ω_2. For example, if $\omega_1 = 0\%$ and $\omega_2 = 2\%$, the bet is always over a 2-percentage-point range over the index value as of $t+1$ regardless of what the value of σ_t^2 is. The expected return therefore depends only on the price of risk γ_t.

Equation (38) characterizes the one-period (or short-horizon) log expected return on the above option strategy. We now consider the two-period (long-horizon) return on this strategy, in order to characterize the term structure of expected returns. To maintain notation, continue to set the option expiration date to $T = t+1$, but now step back to period $t-1$, two periods from expiration. The strategy's log expected return as of tomorrow (date t) will, as in (38), be higher given a positive shock to γ_t. If $\text{Cov}_{t-1}(\gamma_t, M_t) < 0$ — that is, if the price per unit of risk is higher in good times, as we found empirical evidence for in the previous section — then this implies that the return on the strategy from $t-1$ to t will be positive in bad times.[12] The strategy thus provides a hedge against shocks to M_t, implying that its two-period expected return should be lower than its one-period expected return.[13] This implies a downward-sloping term structure of risk prices for option portfolios that fix

[12]Note that the unexpected return on the strategy from $t-1$ to t depends on not just the expected return from t to $t+1$, but also on $\log f_t(\omega_2) - \log f_t(\omega_1)$: the unexpected log return, from (36), depends on $\log f_t^*(\omega_2) - \log f_t^*(\omega_2) = \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_2] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_1] + \log f_t(\omega_2) - \log f_t(\omega_1)$. One concern might be that $\log f_t(\omega_2) - \log f_t(\omega_1)$ decreases in bad times enough to in fact make the strategy have a negative unexpected return in these times. But as shown in Appendix B.1, one can guarantee that $\log f_t(\omega_2) - \log f_t(\omega_1)$ increases in bad times (thereby guaranteeing that the unexpected return is higher in these times) by focusing on sufficiently high return states ω_1 and ω_2. (And more generally, the change in $\log f_t(\omega_2) - \log f_t(\omega_1)$ is likely to be quite small in practice.) So this is not, in our view, a first-order concern.

[13]A full formal analysis of the two-period expected return would require fully specifying the dynamics of all the state variables γ_t, σ_t^2, and so on.

the quantity of risk in the sense described above. We now turn to index options data for suggestive empirical evidence along these lines.

Evidence[14]

We now seek to estimate returns on the fixed-quantity-of-risk strategy $\mathbb{E}_t[R_{\omega,T}]$, varying the horizon $T-t$. Our test of interest is whether this average return decreases in the horizon $T-t$, which (following the discussion above) would provide indirect ex ante evidence that γ_t is higher in good times.

As in the first line of (38), estimating $\mathbb{E}_t[R_{\omega,t+1}]$ is equivalent to estimating the ratio $\frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}$. This ratio of SDF realizations across return states, holding fixed the difference in returns, can be estimated using risk-neutral probabilities $f_t^*(\omega) \equiv f_t^*(R_{m,0\to T}=\omega)$ obtained from index option prices. In particular, we use the same Breeden and Litzenberger (1978)–based approach as described in Section 3 to back out a discretized distribution $f_t^*(\omega)$ across possible returns ω realized over the life of the option. We then translate these into a set of conditional probabilities over binary outcomes — in particular, the probability that the index return from 0 to T will be ω_1 conditional on it being either ω_1 or ω_2. To estimate $\frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}$, we then use the fact that

$$\frac{\pi_t^*}{1-\pi_t^*} = \phi_{t,T}\,\frac{\pi_t}{1-\pi_t}, \tag{39}$$

$$\text{where}\quad \pi_t^* \equiv f_t^*(R_{m,0\to T}=\omega_1 \mid R_{m,0\to T}\in\{\omega_1,\omega_2\}),$$

$$\pi_t \equiv f_t(R_{m,0\to T}=\omega_1 \mid R_{m,0\to T}\in\{\omega_1,\omega_2\}),$$

$$\phi_{t,T} \equiv \frac{\mathbb{E}_t[M_T|R_{m,T}=\omega_1]}{\mathbb{E}_t[M_T|R_{m,T}=\omega_2]}.$$

While we can measure π_t^* from index options data directly, we must estimate π_t across horizons by using the fact that it must be an unbiased forecast of the terminal outcome $\mathbb{1}(R_{m,0\to T}=\omega_1 \mid R_{m,0\to T}\in\{\omega_1,\omega_2\})$ by definition. That is, we are effectively estimating the price of risk embedded in the π_t^* values across horizon such that the implied π_t values have zero average forecast error for the terminal index outcome. We provide formal details of this approach in the appendix (see Appendix B.2).

For implementation, we use S&P 500 index options data from the OptionMetrics database for the period 1996–2018. This yields data for 5,537 trading dates and 991

[14] The empirical results described here were originally reported in Lazarus (2022), which this paper now supersedes.

expiration dates. We drop any options with bid prices of zero (or less than zero), with Black-Scholes implied volatility of greater than 100 percent, or with greater than 12 weeks to maturity (given the relative lack of observations and statistical power beyond this maturity), and calculate each option's end-of-day price as the midpoint between its bid and ask prices.

For each observed expiration date T and associated initial option trading date 0, we define the relevant (sub)set of possible terminal index returns as

$$\Omega = R^f_{0,T} \exp\Big(\{[-0.10, -0.08), [-0.08, -0.06), \ldots, [0.06, 0.08), [0.08, 0.10)\}\Big).$$

In words, state ω_1 is said to be realized when the gross index-price appreciation, in excess of the risk-free rate $R^f_{0,T}$, is between $\exp(-0.1)$ and $\exp(-0.08)$, or equivalently when the log excess return is between -10% and -8%, and analogously for ω_2, and so on. We exclude all terminal states more than 10% out of the money (where moneyness is relative to a zero excess return) in either direction. Note that the states are equally spaced, and all binary bets (e.g., ω_2 vs. ω_3, or ω_5 vs. ω_6) have the same fixed 2-percentage point range of return outcomes within a given option contract, as required by construction. For a given option contract (i.e., a given set of option prices observed from 0 to T), we consider only (ω_i, ω_{i+1}) pairs for which the realized index return was either ω_i or ω_{i+1}. (Without this conditioning, the conditional physical probabilities would be undefined.) This leaves 549 observations (tuples (t, T, i)) at the one-day horizon, which declines monotonically to 222 observations at the 60-day horizon (equivalently, the 12-week horizon), which motivates our focus on 1- to 12-week horizons.

We present the option-implied prices of risk by horizon $\kappa = T - t$ in Figure 3. As the graph shows, the estimated price of risk is significantly downward-sloping as one increases the horizon κ. In other words, one needs a lower price of risk to rationalize the returns on fixed-quantity-of-risk bets at longer horizons. Equivalently, these bets have lower expected returns at longer horizons. As discussed at the end of the previous subsection, this implies that $\text{Cov}_{t-1}(\gamma_t, M_t) < 0$, so that the price per unit of risk is higher given good shocks. This is required in order for the strategy to provide a hedge against shocks to M_t and have lower expected returns at longer horizons.

These results are, in effect, an out-of-sample test in support of the procyclicality of γ_t established earlier in the paper. As discussed theoretically above, the option-implied term structure considered here is downward-sloping if and only if γ_t is procyclical ex ante. The fact that we indeed find a downward-sloping term structure therefore provides further support to the preceding results based on ex post returns on the market.

That said, due to the relatively short horizon necessitated by our data sample and cleaning, the evidence obtained from this estimation is at most suggestive: there is a downward-sloping term structure of implied expected returns on the above option strategy over a matter of weeks, implying that γ_t is procyclical at least at a weekly to monthly data frequency. Whether this speaks to slightly lower frequencies of data aggregation remains an interesting question for future work.

5.3. Additional implications for the equity term structure

The preceding discussion considered option-implied risk prices by horizon, which constitute a particular equity term structure in which one holds fixed the riskiness of the claim in question and varies the horizon. This is slightly different than the notion of the equity term structure studied in much of the recent literature (e.g., van Binsbergen, Brandt, and Koijen, 2012), which considers index dividends at different horizons and studies their (risk-adjusted) returns. While the two concepts are somewhat distinct, all our findings (for both equity returns and option-implied risk prices) can be reconciled with stylized facts about the standard equity term structure. We consider a model building off of that of Lettau and Wachter (2007), but with time-varying volatility in both discount-rate shocks and dividend-growth shocks. We will see that a very stripped down model of this form will produce a generally procyclical price of variance risk but a countercyclical Sharpe ratio, and it will also predict a downward-sloping term structure of risk-adjusted returns.[15]

The aggregate dividend is denoted by D_t, and let $d_t = \log D_t$. We assume that log dividend growth follows

$$\Delta d_{t+1} = g - \frac{1}{2}x_t^2\sigma_z^2 + \sigma_d\varepsilon_{d,t+1} + x_t\sigma_z\varepsilon_{z,t+1}, \tag{40}$$

where $\varepsilon_{d,t+1}$ and $\varepsilon_{z,t+1}$ are standard normal and independent of each other and over time. Relative to the specification in Lettau and Wachter (2007), we include an additional shock $x_t\sigma_z\varepsilon_{z,t+1}$, whose volatility is time-varying and increasing in x_t. This x_t variable will also, for parsimony, represent the price of risk. The shock $\varepsilon_{d,t+1}$ will be priced (i.e., it will enter the SDF), while the shock with time-varying volatility will not be. This should be thought of as a stripped-down way to model the idea that returns on dividend strips include additional "non-fundamental" volatility in bad times, reverse-engineered here

[15] Because such a model features a downward-sloping term structure of risk-adjusted returns in spite of the time variation in volatility, it also by implication would feature a downward-sloping term structure for strategies that hold fixed the quantity of risk, thereby reconciling with the results in the previous section. And if, in addition, we included a small price of risk on discount-rate shocks, one could reconcile this model with the countercyclical term premium as in Gormsen (2021).

by including an unpriced shock in dividend growth that becomes more important in bad times.[16]

As above, the price of risk is driven by a single state variable x_t, which follows

$$x_{t+1} = (1-\phi_x)\overline{x} + \phi_x x_t + \sqrt{x_t}\sigma_x \varepsilon_{x,t+1}, \tag{41}$$

where $\varepsilon_{x,t+1}$ is i.i.d. standard normal and independent of $\varepsilon_{d,t+1}$ and $\varepsilon_{z,t+1}$. The time-varying volatility of the price of risk, governed by the square-root process implied by (41), is a further distinction from Lettau and Wachter (2007). Assume that $\overline{x} > 0$ and $\phi_x \in (0,1)$, so that x_t is positive with probability one in a continuous-time limit of the model. We also assume that $\sigma_z \overline{x} \geqslant \sigma_d$ so that in (40), "non-fundamental" dividend shocks are at least as important as "fundamental" dividend shocks in steady state.

The log stochastic discount factor $m_{t+1} = \log M_{t+1}$ is directly specified as

$$m_{t+1} = -r^f - \frac{1}{2}x_t^2 - x_t \varepsilon_{d,t+1}, \tag{42}$$

where r^f is the constant log risk-free rate. Intuitively, investors dislike exposure to "fundamental" dividend-growth shocks $\varepsilon_{d,t+1}$, and the degree to which they dislike this exposure is governed by risk aversion (the conditional price of risk) x_t. All other risks are unpriced directly. This is an extreme assumption, but again it clarifies economic intuition.

We can solve explicitly for the prices and returns of zero-coupon equity (i.e., n-maturity dividend claims).[17] The price of the n-maturity claim at time t is $P_{n,t}$, and let $p_{n,t} = \log P_{n,t}$. One-period returns are $R_{n,t+1} = P_{n-1,t+1}/P_{n,t}$. Since $\mathbb{E}_t[M_{t+1}R_{n,t+1}] = 1$, we have the following recursive relation for prices:

$$P_{n,t} = \mathbb{E}_t[M_{t+1}P_{n-1,t+1}], \tag{43}$$

with boundary condition $P_{0,t} = D_t$ given that the dividend is paid out at maturity.

Guess a log-linear solution for the price-dividend ratio:

$$\frac{P_{n,t}}{D_t} = \exp(A_n + B_{x,n}x_t)\,. \tag{44}$$

[16] As an additional twist relative to Lettau and Wachter (2007), we also rule out time variation in the conditional mean of (exponentiated) dividend growth. Including such time variation would allow for a downward-sloping term structure of expected returns (rather than a constant term structure of expected returns but downward-sloping Sharpe ratios and CAPM alphas, as ours will feature). Since this is relatively unimportant for our analysis, we simplify by omitting such time variation.

[17] The price and return for aggregate equity then follows straightforwardly from the zero-coupon solutions, but in order to examine intuition, we maintain focus on the zero-coupon claims.

Under this conjecture, the price-dividend ratio is

$$\frac{P_{n,t}}{D_t} = \mathbb{E}_t \left[M_{t+1} \frac{D_{t+1}}{D_t} \exp(A_{n-1} + B_{x,n-1} x_{t+1}) \right]. \tag{45}$$

Using the assumed conditional log-normality in (40)–(41), we can match coefficients of (44) and (45) to obtain

$$A_n = A_{n-1} - r^f + g + B_{x,n-1}(1 - \phi_x)\overline{x} + \frac{1}{2}\sigma_d^2, \tag{46}$$

$$B_{x,n} = B_{x,n-1}\left(\phi_x + \frac{1}{2}B_{x,n-1}\sigma_x^2\right) - \sigma_d, \tag{47}$$

with boundary conditions $A_0 = B_{x,0} = 0$. This verifies the conjecture. Note that $B_{x,n} < 0$ for all n under the weak condition that $\sigma_d \sigma_x^2 < 2\phi$, so that the price-dividend ratio decreases (times are bad) when the price of risk increases.

The log return on the strip of maturity n is thus

$$\begin{aligned} r_{n,t+1} &= \log\left(\frac{P_{n-1,t+1}}{D_{t+1}} \frac{D_t}{P_{n,t}} \frac{D_{t+1}}{D_t}\right) \\ &= g - \frac{1}{2}x_t^2\sigma_z^2 + \sigma_d \varepsilon_{d,t+1} + x_t \sigma_z \varepsilon_{z,t+1} + A_{n-1} + B_{x,n-1}x_{t+1} - A_n - B_{x,n}x_t. \end{aligned} \tag{48}$$

The conditional variance of this log return follows as

$$\sigma_{n,t}^2 = \text{Var}_t(r_{n,t+1}) = \sigma_d^2 + \sigma_z^2 x_t^2 + |B_{x,n-1}|^2 \sigma_x^2 x_t. \tag{49}$$

The excess expected return, meanwhile, is

$$\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2 = -\text{Cov}_t(r_{n,t+1}, m_{t+1}) = \sigma_d x_t, \tag{50}$$

so the term structure of expected returns is flat.

Putting (49) and (50) together, the Sharpe ratio is

$$\text{SR}_{n,t} \equiv \frac{\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2}{\sigma_{n,t}} = \frac{\sigma_d x_t}{\sqrt{\sigma_d^2 + \sigma_z^2 x_t^2 + |B_{x,n-1}|^2 \sigma_x^2 x_t}}. \tag{51}$$

Note first that this is decreasing in maturity n, so we obtain a downward-sloping term structure of Sharpe ratios. In addition, the Sharpe ratio is countercyclical:

$$\frac{\partial \mathrm{SR}_{n,t}}{\partial x_t} \propto \left(\sigma_d^2 + \frac{1}{2} B_{x,n-1}^2 \sigma_x^2 x_t \right) > 0. \tag{52}$$

The countercyclical price of risk passes through to generate a countercyclical Sharpe ratio, as is standard.

But the ratio of expected returns to *variance*, meanwhile, is

$$\gamma_{n,t} \equiv \frac{\mathbb{E}_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_{n,t}^2}{\sigma_{n,t}^2} = \frac{\sigma_d x_t}{\sigma_d^2 + \sigma_z^2 x_t^2 + |B_{x,n-1}|^2 \sigma_x^2 x_t}, \tag{53}$$

which varies with x_t according to

$$\frac{\partial \gamma_{n,t}}{\partial x_t} \propto \left(\sigma_d^2 - \sigma_z^2 x_t^2 \right). \tag{54}$$

This value can be either positive or negative, and it will be negative if and only if $\sigma_z x_t > \sigma_d$. So for x_t large enough, we obtain a *procyclical* price per unit of variance risk: further positive shocks to x_t increase non-fundamental return volatility enough to offset the increase in expected returns. Under the maintained assumption that $\sigma_z \overline{x} > \sigma_d$, this procyclicality holds around steady state, though this is not particularly important for our analysis. We predict in general that there will be a countercyclical $\gamma_{n,t}$ for a part of the state space (in good times, when x_t is small), as we observe in the time-series data. But in bad enough times, non-fundamental return volatility is large enough to make the price per unit of variance risk decrease given further increases in x_t.

Intuitively, expected returns go up with the price of risk, but the importance of non-fundamental risk for returns also increases. While return volatility increases, the beta of the SDF onto the market decreases during these times because of the rise in non-fundamental return risk. In other words, return volatility has three components: fundamental dividend volatility (which is constant, σ_d), discount-rate volatility (which increases in $\sqrt{x_t}$), and non-fundamental volatility (which increases in x_t). An increase in x_t therefore increases "pure" market risk that is not fully connected to fundamentals, increasing market variance without passing through one-for-one to expected returns. While this is sufficient to generate a procyclical $\gamma_{n,t}$ that decreases in x_t (at least for large x_t), this non-fundamental volatility effect is not strong enough to obtain a procyclical Sharpe ratio. Further, the greater exposure of long-maturity claims to discount-rate risk means that their volatility increases without changing their expected return, generating a downward-sloping Sharpe ratio of dividend claims.

As a result, this stylized model shows how our findings about the cyclicality of $\gamma_{n,t}$ connect to facts about the dividend term structure. We obtain *both* (i) a lower beta of the SDF onto the market in bad times (our main stylized fact), and (ii) a downward-sloping Sharpe ratio of dividend claims by maturity, both through the same channel (non-fundamental return risk). Meanwhile, our parameterization maintains the usual countercyclical Sharpe ratio.

With respect to the *cyclicality* of the term structure, speaking to the countercyclicality documented by Gormsen (2021) could be achieved by assuming, similar to Gormsen, that the discount-rate shock $\varepsilon_{x,t+1}$ also enters into the SDF, with a small average price of risk for this shock but an increase in the quantity of this risk in bad times.[18]

6. Conclusion

We show that there is a reverse cyclicality in two seemingly similar risk prices, the price per unit of volatility risk, the Sharpe ratio of market market returns, and the price per unit of variance risk. In a global sample covering 20 stock markets around the world, we show that the Sharpe ratio is countercyclical, consistent with conventional wisdom, but the price per unit of variance risk is procyclical. This implies that every unit of return variance matters less for investors in bad times than in good times, as we show theoretically. The procyclicality of the price per unit of variance risk has implications for the term structure of expected returns on constant risk option portfolios and for the equity term structure more generally. We provide a theoretical link between the cyclicality in the price per unit of variance risk and the equity term structure and provide empirical evidence in favor of our theory.

[18]Such a model would also allow us to speak to the preceding subsection's results on risk prices for strategies that hold fixed the quantity of risk. That analysis required that $\mathrm{Cov}_t(\gamma_{t+1}, M_{t+1}) < 0$, but in the current simple case, $\gamma_{d,t}$ in (53) has no covariance with the SDF. Adding an x_{t+1} shock to the SDF would accordingly also accommodate the previous analysis and allow for a more formal tie between the two sets of results.

References

AIT-SAHALIA, Y. AND A. W. LO (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94, 9—-51.

BATES, D. S. (2000): "Post 87 Crash Fears in the S&P 500 Futures Options Market," *Journal of Econometrics*, 94, 181–238.

BERKOWITZ, J. (2001): "Testing Density Forecats, With Applications to Risk Management," *Journal of Business and Economic Statistics*, 19, 465–474.

VAN BINSBERGEN, J., M. BRANDT, AND R. KOIJEN (2012): "On the timing and pricing of dividends," *American Economic Review*, 102, 1596–1618.

BLISS, R. R. AND N. PANIGIRTZOGLOU (2004): "Option Implied Risk Aversion Estimates," *Journal of Finance*, 1, 407–446.

BREEDEN, D. T. AND R. H. LITZENBERGER (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *The Journal of Business*, 51, 621—-651.

CAMPBELL, J. Y. (2018): *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton University Press.

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.

CAMPBELL, J. Y., S. GIGLIO, R. TURLEY, AND C. POLK (2018): "An intertemporal CAPM with stochastic volatility," *Journal of Financial Economics*, 128, 207–233.

CHAUVET, M. AND J. PIGER (2008): "A comparison of the real-time performance of business cycle dating methods," *Journal of Business and Economic Statistics*, 26, 42–49.

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

FAMA, E. F. (1968): "Risk, Return and Equilibrium: Some Clarifying Comments," *Journal of Finance*, 23, 29–40.

FIGLEWSKI, S. (2018): "Risk Neutral Densities: A Review," *working paper*.

FRIEND, I. AND M. E. BLUME (1975): "The Demand for Risky Assets," *American Economic Review*, 65, 900–922.

GORMSEN, N. J. (2021): "Time Variation of the Equity Term Structure," *Journal of Finance*, 76, 1959–1999.

GORMSEN, N. J. AND C. S. JENSEN (2024): "Higher Moment Risk," *Journal of Finance, forthcoming*.

JACKWERTH, J. (2000): "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

JACKWERTH, J. C. (2004): "Option-Implied Risk-Neutral Distributions and Risk Aversion," *CFA Institute Research Foundation of AIMR Publications*, 1–86.

JENSEN, C. S. (2018): "Essays on Asset Pricing," *PhD dissertation, Copenhagen Business School*.

JENSEN, C. S., D. LANDO, AND L. H. PEDERSEN (2019): "Generalized Recovery," *Journal of Financial Economics*, 133(1), 154–174.

KELLY, B. AND S. PRUITT (2013): "Market Expectations in the Cross-Section of Expected Values," *Journal of Finance*, 68, 1721–1756.

LAZARUS, E. (2022): "Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options," *Working paper*.

LETTAU, M. AND J. A. WACHTER (2007): "Why is long-horizon equity less risky? a duration-based explanation of the value premium," *The Journal of Finance*, 62, 55–92.

LUSTIG, H. AND A. VERDELHAN (2012): "Business cycle variation in the rreturn trade-off," *Journal of Monetary Economics*, 57, 535–549.

MOREIRA, A. AND T. MUIR (2017): "Volatility-Managed Portfolios," *Journal of Finance*, 72, 1611—-1644.

ROSENBERG, J. V. AND R. F. ENGLE (2002): "Empirical Pricing Kernels," *Journal of Financial Economics*, 64, 341—-372.

SCHREINDORFER, D. AND T. SICHERT (2023): "Volatility and the Pricing Kernel," *working paper*.

Figure 1: **Time variation in risk prices**. Subfigure (a) shows the Sharpe ratio of the S&P 500 index. Subfigure (b) shows the price per unit of variance risk. The grey shaded areas are NBER recession periods. We standardize the measures to make them easily comparable in the figure. Shaded area is NBER recession periods.

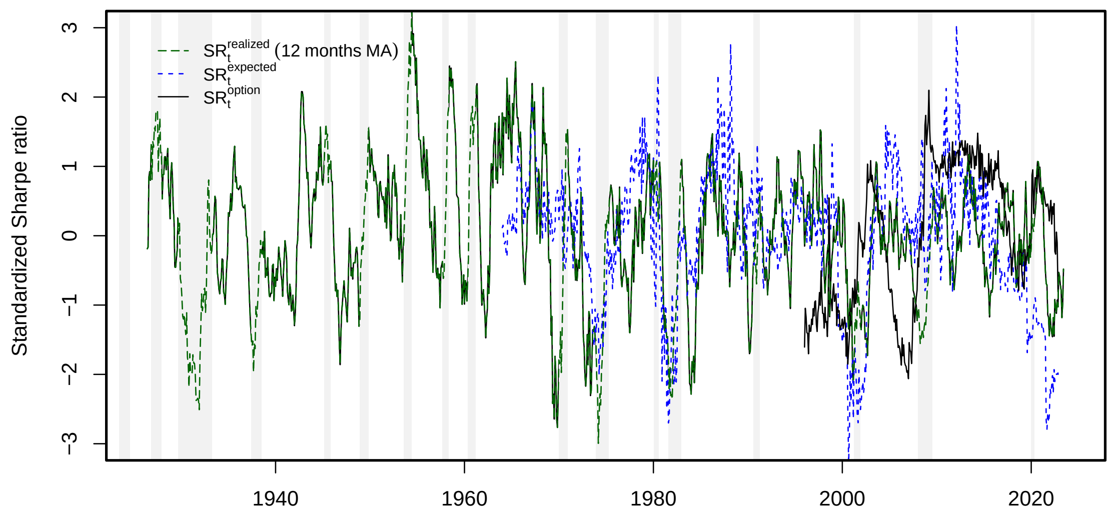

(a) Sharpe ratio

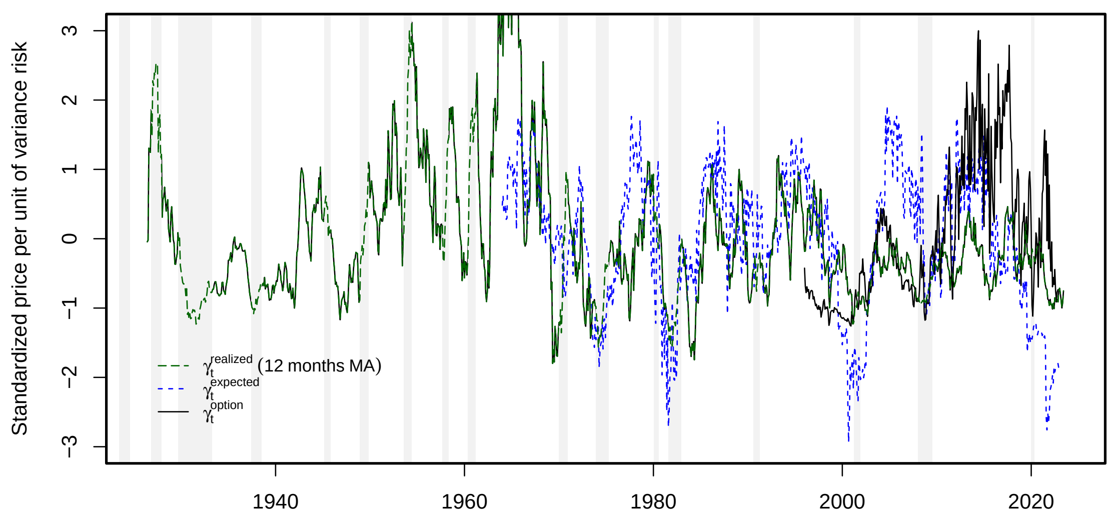

(b) Price per unit of variance risk

Figure 2: **Prices of risk**. This figure shows the option implied price per unit of variance risk and the option implied Sharpe ratio for the S&P 500 index. The measures are standardized to make them easily comparable in the figure. Shaded aread are NBER recession periods.

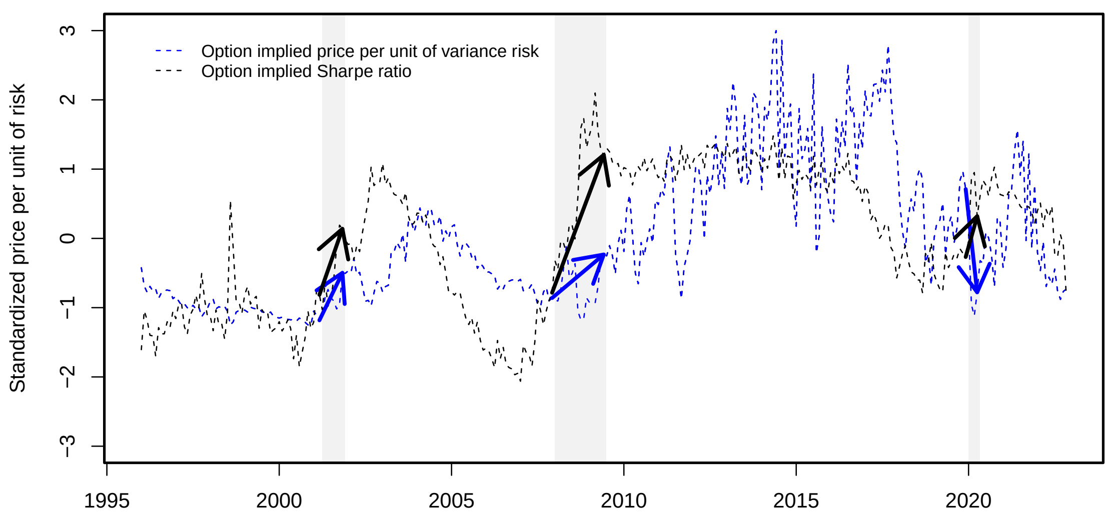

Figure 3: **Estimates of Risk Prices by Horizon**. This figure shows the option-implied price of risk by horizon. Point estimates are constructed using two-step GMM, using the five-day-lagged observation as an instrument as described in Appendix B.2, on the sample counterparts of the moment conditions described in the appendix in order to minimize forecast error. The price of risk parameter is constrained to be equal for all days within a given weekly horizon to expiration. Error bars show 95% confidence intervals, constructed using procedure in Appendix B.2.

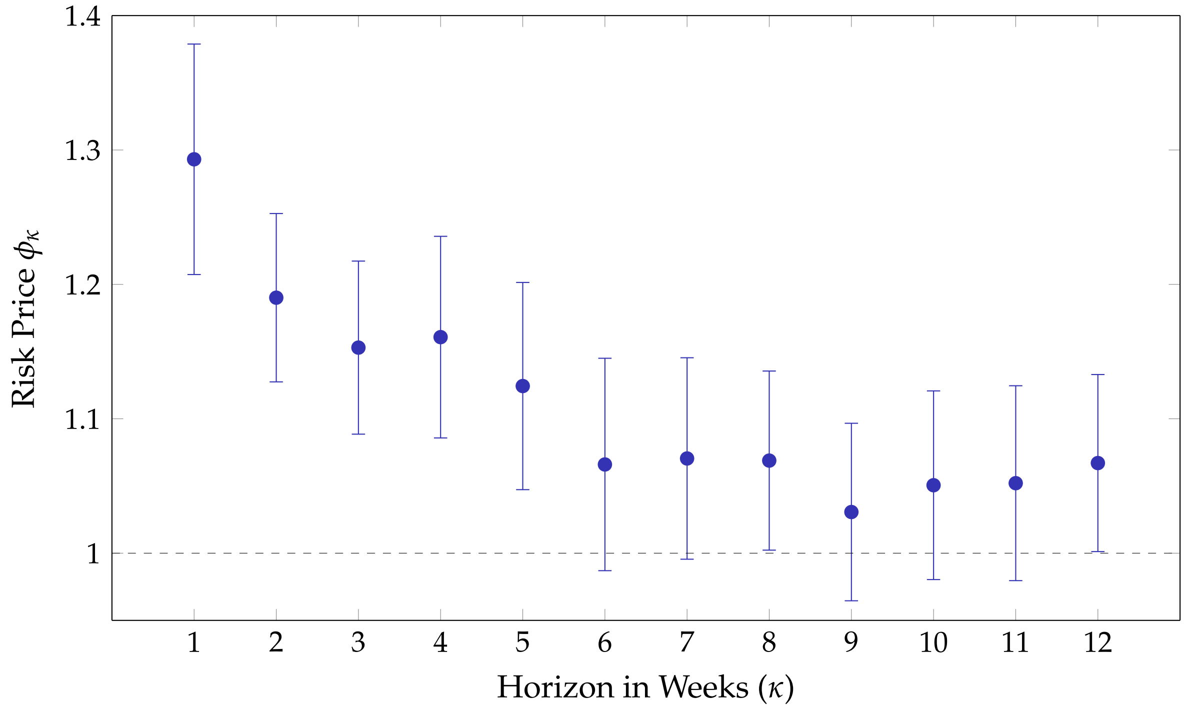

Table 1: **Pairwise correlations of risk prices.**
This table reports the pairwise correlations of the price per unit of variance risk (γ_t) and the Sharpe ratio (SR_t) for the S&P 500 index. We compute four measures for each of the risk prices: (i) a 12-months moving average of the monthly realized measure $\gamma_t^{\text{realized}}$, using daily returns to compute conditional expected excess returns and volatilities, (ii) $\gamma_t^{\text{expected}}$, as an expected measure using Kelly and Pruitt (2013) for expected excess returns and an AR(1) process for the variance, (iii) γ_t^{option}, an option implied measure estimated from options written on the S&P 500 index, and (iv) γ_t^{PCA}, the first principal component of the three previous measures.

	$\gamma_t^{\text{expected}}$	γ_t^{option}	γ_t^{PCA}	SR_t^{realized}	SR_t^{expected}	SR_t^{option}	SR_t^{PCA}
$\gamma_t^{\text{realized}}$	0.42	0.29	0.46	0.84	0.30	−0.09	0.28
$\gamma_t^{\text{expected}}$		0.27	0.66	0.43	0.91	−0.03	0.66
γ_t^{option}			0.78	0.28	0.22	0.54	0.51
γ_t^{PCA}				0.42	0.59	0.25	0.77
SR_t^{realized}					0.39	0.08	0.41
SR_t^{expected}						0.13	0.80
SR_t^{option}							0.54

Table 2: **Pre-recession to end-of-recession changes in prices of risk.**
This table reports the average changes in the risk prices of the S&P 500 index from the month before the onset of a recession to the end of the last month in the recession. We standardize the risk measures to make them comparable. A coefficient of 1 should be interpreted as a one unconditional standard deviation increase in the risk price. The annualized standard deviation of the Sharpe ratio is about 0.17 in both the expected and option implied samples. The average annualized Sharpe ratio is about 0.48 for both the expected and option implied samples. Statistical significance at the 10% level is shown in bold.

	Realized	Expected	Option	PCA
Δ Sharpe ratio (ΔSR_t)	**1.28**	**0.58**	1.18	**1.28**
t-stat	4.44	2.46	2.80	3.19
Δ Price per unit of variance risk ($\Delta\gamma_t$)	**0.77**	0.09	−0.06	0.04
t-stat	2.98	0.51	−0.08	0.05
ΔSR_t - $\Delta\gamma_t$	**0.51**	**0.49**	1.24	**1.24**
t-stat	3.23	3.27	2.35	3.55
No. recessions	16	8	3	3

Table 3: **Cyclicality in the Sharpe Ratio of market returns — With perfect foresight in recessions.**

This table compares the Sharpe ratio for the US stock market in normal times (NBER non-recession months) to the months in recessions **after** the stock market reached its low during the recession. Specifically, we compare the prices of risk for an investor who only invests in good times to that of an investor who has perfect foresight during recessions in the sense that she can pinpoint when the market has reached its low. This investors buys the market at its low and holds the market for twelve months. We compute unconditional measures by bundling monthly returns based on each trading strategy and compute within bundle expected excess returns and variance. The "Difference" rows report the difference between recession periods and normal times. Statistical significance at the 10% level is shown in bold. *t*-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors.

	Full sample	Post 1964 sample	Post 1996 sample
Unconditional Sharpe Ratio			
In recessions (after stock market low)	0.51	0.59	0.49
In normal times	0.17	0.15	0.20
Difference (recession - normal)	**0.34**	**0.44**	0.29
Standard error	0.10	0.16	0.27
Conditional realized Sharpe Ratio			
In recessions (after stock market low)	3.39	2.62	1.85
In normal times	1.61	1.49	1.42
Difference (recession - normal)	**1.78**	**1.13**	0.43
t-stat	4.45	2.19	0.52
Conditional expected Sharpe Ratio			
In recessions (after stock market low)		0.36	0.20
In normal times		0.53	0.46
Difference (recession - normal)		$-$**0.17**	-0.26
t-stat		-1.69	-1.15
Conditional option implied Sharpe Ratio			
In recessions (after stock market low)			0.55
In normal times			0.45
Difference (recession - normal)			**0.10**
t-stat			1.91
Summary statistics			
No. recession periods	16	8	3
No. recession months after market low	192	96	36
No. normal times months	865	556	277

Table 4: **Cyclicality in the price per unit of variance risk — With perfect foresight in recessions.**

This table compares the price per unit of variance risk for the US stock market in normal times (NBER non-recession months) to the months in recessions **after** the stock market reached its low during the recession. Specifically, we compare the prices of risk for an investor who only invests in good times to that of an investor who has perfect foresight during recessions in the sense that she can pinpoint when the market has reached its low. This investors buys the market at its low and holds the market for twelve months. We compute unconditional measures by bundling monthly returns based on each trading strategy and compute within bundle expected excess returns and variance. The ”Difference” rows report the difference between recession periods and normal times. Statistical significance at the 10% level is shown in bold. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors.

	Full sample	Post 1964 sample	Post 1996 sample
Unconditional price per unit of variance risk			
In recessions (after stock market low)	7.69	10.77	7.55
In normal times	3.62	3.27	4.12
Difference (recession - normal)	**4.07**	7.50	3.44
Standard error	2.44	3.76	5.95
Conditional realized price per unit of variance risk			
In recessions (after stock market low)	40.40	21.53	11.20
In normal times	26.04	23.32	15.98
Difference (recession - normal)	**14.36**	−1.79	−4.78
t-stat	1.75	−0.31	−0.96
Conditional expected price per unit of variance risk			
In recessions (after stock market low)		1.59	0.72
In normal times		3.35	2.95
Difference (recession - normal)		**−1.77**	**−2.23**
t-stat		−3.65	−2.46
Conditional option implied price per unit of variance risk			
In recessions (after stock market low)			2.54
In normal times			3.20
Difference (recession - normal)			**−0.65**
t-stat			−2.06
Summary statistics			
No. recession periods	16	8	3
No. recession months after market low	192	96	36
No. normal times months	865	556	277

Table 5: **Cyclicality in the Sharpe Ratio of market returns — Without perfect foresight in recessions.**
This table compares the Sharpe ratio for the US stock market in normal times (NBER non-recession months) to what is earned in recessions for an investor who buys the market six months into each recession and holds the market for twelve months thereafter. We focus exclusively on recessions where the sixth month after the onset of a recession is still a recession month. We compute unconditional measures by bundling monthly returns based on each trading strategy and compute within bundle expected excess returns and variance. The "Difference" rows report the difference between recession periods and normal times. Statistical significance at the 10% level is shown in bold. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors.

	Full sample	Post 1964 sample	Post 1996 sample
Unconditional Sharpe Ratio			
In recessions (after stock market low)	0.23	0.12	−0.16
In normal times	0.18	0.18	0.24
Difference (recession - normal)	0.04	−0.06	−0.39
Standard error	0.11	0.16	0.32
Conditional realized Sharpe Ratio			
In recessions (after stock market low)	2.40	1.13	−0.09
In normal times	1.63	1.55	1.51
Difference (recession - normal)	0.77	−0.42	**−1.60**
t-stat	1.31	−0.64	−1.90
Conditional expected Sharpe Ratio			
In recessions (after stock market low)		0.27	0.17
In normal times		0.53	0.45
Difference (recession - normal)		**−0.25**	−0.28
t-stat		−2.79	−1.48
Conditional option implied Sharpe Ratio			
In recessions (after stock market low)			0.55
In normal times			0.46
Difference (recession - normal)			0.09
t-stat			1.13
Summary statistics			
No. recession periods	16	8	3
No. recession months after market low	180	84	24
No. normal times months	877	578	296

Table 6: **Cyclicality in the price per unit of variance risk — Without perfect foresight in recessions.**
This table compares the price per unit of variance risk for the US stock market in normal times (NBER non-recession months) to what is earned in recessions for an investor who buys the market six months into each recession and holds the market for twelve months thereafter. We focus exclusively on recessions where the sixth month after the onset of a recession is still a recession month. We compute unconditional measures by bundling monthly returns based on each trading strategy and compute within bundle expected excess returns and variance. The "Difference" rows report the difference between recession periods and normal times. Statistical significance at the 10% level is shown in bold. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors.

	Full sample	Post 1964 sample	Post 1996 sample
Unconditional price per unit of variance risk			
In recessions (after stock market low)	3.63	1.82	−1.83
In normal times	3.73	3.74	4.90
Difference (recession - normal)	−0.09	−1.92	−6.72
Standard error	2.33	3.35	6.63
Conditional realized price per unit of variance risk			
In recessions (after stock market low)	33.81	10.99	1.03
In normal times	25.68	23.33	16.14
Difference (recession - normal)	8.14	**−12.34**	**−15.12**
t-stat	0.87	−1.70	−3.62
Conditional expected price per unit of variance risk			
In recessions (after stock market low)		1.24	0.68
In normal times		3.29	2.83
Difference (recession - normal)		**−2.05**	**−2.16**
t-stat		−4.29	−2.61
Conditional option implied price per unit of variance risk			
In recessions (after stock market low)			2.03
In normal times			3.18
Difference (recession - normal)			**−1.15**
t-stat			−3.54
Summary statistics			
No. recession periods	16	8	3
No. recession months after market low	180	84	24
No. normal times months	865	556	277

Table 7: **Cyclicality in the price of risk - Pooled international evidence.**
For 20 stock market indexes, this table reports the results of pooled panel regressions of the differences in the conditional monthly price per unit of variance risk during normal times (OECD non-recession months) to that of an investor who invests for twelve months during recessions, starting either: (i) **after** the stock market reached its low during the recession, (ii) one month into each recession, (iii) six months into each recession, or (iv) twelve months into each recession. We report differences as recession − normal times, a positive value means that the price of risk is higher in recessions. We first compute conditional prices of risk within each month and thereafter investigate the average conditional prices of risk in normal times versus in recessions. We compute the price per unit of variance risk in two ways: (i) a realized measure $\gamma_t^{\text{realized}}$, using daily returns to compute conditional expected excess returns and volatilities and (ii) $\gamma_t^{\text{expected}}$, as an expected measure using Kelly and Pruitt (2013) for expected excess returns and an AR(1) process for the variance. Statistical significance at the 5% level is shown in bold. We include stock index fixed effects and cluster standard errors by index and date.

	Sharpe ratio	$\gamma_t^{\text{realized}}$	$\gamma_t^{\text{expected}}$
Market low in recession	**0.17**	0.99	−0.41
t-stat	2.48	0.65	−1.26
One month into recession	**−0.33**	**−7.45**	**−1.05**
t-stat	−5.04	−5.07	−3.53
Six months into recession	**−0.29**	**−6.80**	**−1.15**
t-stat	−4.64	−5.14	−3.96
Twelve months into recession	**−0.15**	**−3.61**	**−0.91**
t-stat	−2.35	−2.45	−3.21

Table 8: **Consumption growth, financial and macroeconomic conditions, and the price of risk.**
This table reports the results of regressions:

$$\text{Price of risk}_t = \alpha + \beta \times \text{Indicator}_t + \epsilon_t \qquad (55)$$

t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. Data on financial risk indicator is at the monthly horizon and obtained from the Chicago fed database. NFCI is the national financial condition indicator. According to the Chicago fed, "Risk" captures volatility and funding risk in the financial sector. Credit captures credit conditions and leverage consists of debt and equity measures. High values of the variables are historically associated with tighter-than-average conditions in financial markets, i.e., bad times. CFNAI is the Chicago Fed National Activity Index, build to capture movements in economic expansions and contractions as well as periods of increasing and decreasing inflationary pressure. A low value of this variable is typically associated with economic contractions. The "Rec. prob." variable is the recessions probability of Chauvet and Piger (2008). Consumption is the St. Louis Fed monthly growth in personal consumption expenditures of non-durable and service goods, deflated with the CPI. Statistical significance at the 10% level is shown in bold.

	Sharpe ratio			Price per unit of variance risk		
Indicator	$SR_t^{\textbf{realized}}$	$SR_t^{\textbf{expected}}$	$SR_t^{\textbf{option}}$	$\gamma_t^{\textbf{realized}}$	$\gamma_t^{\textbf{expected}}$	$\gamma_t^{\textbf{option}}$
NFCI	**−0.16**	**−0.02**	0.02	**−5.47**	**−0.65**	**−0.88**
t-stat	−2.86	−2.32	1.56	−3.02	−4.03	−2.30
Risk	**−0.16**	**−0.02**	0.02	**−5.43**	**−0.60**	**−0.96**
t-stat	−2.98	−1.82	1.44	−2.98	−3.42	−2.67
Credit	−0.01	−0.02	0.02	−1.49	**−0.49**	**−1.11**
t-stat	−0.16	−1.62	1.05	−0.62	−2.08	−2.28
Leverage	**−0.17**	−0.01	0.01	**−6.21**	−0.34	−0.17
t-stat	−3.05	−0.65	1.04	−3.36	−0.97	−0.35
Non-fin. leverage	**−0.08**	0.01	**−0.02**	−1.78	0.18	**−0.67**
t-stat	−1.82	0.58	−2.29	−1.04	0.70	−4.03
Recession prob.	**−0.46**	**−0.07**	0.03	**−22.71**	**−2.18**	**−1.50**
t-stat	−1.98	−1.89	1.10	−3.49	−2.96	−3.65
CFNAI	0.04	**0.01**	−0.00	3.24	**0.28**	**0.10**
t-stat	0.82	1.65	−0.91	1.52	1.73	1.72
Consumption	**5.24**	0.26	−0.02	352.12	8.41	−7.23
t-stat	1.90	0.64	−0.11	1.56	0.97	−1.35

Table 9: **Further international evidence for macro variables.**
This table reports the results of panel regressions on the form:

$$\text{Risk price}_t^i = \alpha + \beta \times \text{macro variable}_{t+1,t+s}^i + \epsilon_q \tag{56}$$

where i denotes indexes for the up to 20 international stock market indexes in our sample and macro variable$_{t+1,t+s}^i$ is either the: (i) eight quarters growth in consumption, (ii) the contemporaneous value of the market dividend-to-price ratio, or (iii) eight months growth in industrial production. t-statistics are corrected for heteroscedasticity and autocorrelation using Newey West standard errors. Statistical significance at the 10% level is shown in bold. 'Standardized' rows report results where we standardize both the risk price and the macro variable within each country before we pool the data for the regression. 'Raw' rows report results where we pool data without standardizing.

	Sharpe ratio		Price per unit of variance risk	
	SR_t^{realized}	SR_t^{expected}	$\gamma_t^{\text{realized}}$	$\gamma_t^{\text{expected}}$
Consumption growth				
Raw	**0.14**	−0.10	**73.43**	3.35
t-stat	3.61	−0.23	3.29	0.48
Standardized	**0.17**	0.01	**0.14**	0.05
t-stat	4.66	0.14	3.61	0.47
Dividend-price ratio				
Raw	**−4.46**	−0.34	**−99.20**	**−11.28**
t-stat	−2.35	−0.55	−2.07	−1.72
Standardized	**−0.09**	0.03	**−0.08**	−0.04
t-stat	−3.84	0.58	−3.12	−0.81
Industrial production				
Raw	**3.47**	0.22	**58.50**	**5.10**
t-stat	4.67	1.33	3.62	1.83
Standardized	**0.15**	0.07	**0.10**	**0.08**
t-stat	4.43	1.36	3.62	1.78

Appendix

A. Details for Section 3

A.1. Inferring risk-neutral distributions from options

We compute risk-neutral distributions using a spline approach as suggested in Figlewski (2018). On each last trading day of the month, we compute the spline that best fits the observed volatility surface under the two conditions that: (i) the left part of the surface is monotonically decreasing (options with moneyness less than 1) and (ii) the resulting risk-neutral distribution is non-negative. Given a fitted spline, we compute the risk-neutral distribution as the second derivative of the resulting Black-Scholes prices. When we have option data with exactly one month maturity, we simply compute the distribution at this horizon. When we do not have such horizons, we compute distributions for the closest days (from above and below) for which there is data. We then linearly interpolate the distributions for these horizons to obtain a monthly horizon distribution.

There are of course many other ways in which we can compute risk-neutral distributions. For example, we can use a parametric approach like in Bates (2000), the "fast-and-stable" method of Jackwerth (2004), or a number of alternative ways to essentially just smooth implied volatilities to obtain a smooth continuous price function that we can numerically differentiate. Since the Berkowitz (2001) test relies heavily on the first and second moments of the recoverd risk-neutral distributions, which are almost identical using either of these methods, we are confident that alternative methods will yield similar cyclicalities in option implied prices of market variance risk when using the approach described in Section 3.

B. Details for Section 5.2

B.1. Theory

This appendix continues the discussion in footnote 12 on sufficient conditions to guarantee that the unexpected return on the fixed-quantity-of-risk strategy is higher in bad times. As in that footnote, the unexpected log return on the strategy from $t-1$ to t depends on $\log f_t^*(\omega_2) - \log f_t^*(\omega_2) = \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_2] - \log \mathbb{E}_t[M_{t+1}|R_{m,t+1} = \omega_1] + \log f_t(\omega_2) - \log f_t(\omega_1)$ (using (36)). The log-normal density assumption gives that

$$\log f_t(\omega_2) - \log f_t(\omega_1)$$
$$= -\log(\omega_2) + \log(\omega_1) - \frac{\left(\log(\omega_2) - \mu_{R,t} + \frac{1}{2}\sigma_t^2\right)^2 - \left(\log(\omega_1) - \mu_{R,t} + \frac{1}{2}\sigma_t^2\right)^2}{2\sigma_t^2}.$$

This decreases in $\mu_{R,t}$ and it may either increase or decrease in σ_t^2, so one concern (as raised in the footnote) might be that $\log f_t(\omega_2) - \log f_t(\omega_1)$ decreases in bad times enough to in fact make the strategy have a negative unexpected return in these times. But since γ_t decreases in bad times, we must have that $d\gamma_t \propto d\mu_t - \gamma_t d(\sigma_t^2) < 0$ in these times, or $d\mu_t / d(\sigma_t^2) < \gamma_t$. So in order for $\log f_t(\omega_2) - \log f_t(\omega_1)$ to increase in bad times so that the unexpected return is guaranteed to be positive, one can see (after some algebra) that it is sufficient to have

$$\frac{\log(\omega_1) + \log(\omega_2)}{2} \frac{1}{\sigma_t^2} > \gamma_t - d\mu_t / d(\sigma_t^2) > 0.$$

One can always find large enough return states to guarantee that this is the case, meaning that the unexpected return will always be higher in bad times as long as we're focusing on sufficiently high ω_1 and ω_2.

B.2. Empirical evidence: implementation

We begin with equation (39). We are interested in how the price of risk $\phi_{t,T}$ changes on average with the horizon $T - t$, for multiple possible return state pairs (ω_1, ω_2). We therefore assume that for arbitrary pairs of return states (ω_1, ω_2) and (ω_3, ω_4), if $\omega_2/\omega_1 = \omega_4/\omega_3$, then the associated ϕ values are equivalent (i.e., $\phi_{t,T,\omega_1,\omega_2} = \phi_{t,T,\omega_3,\omega_4}$). This is in effect an assumption of scale independence (as would hold under, e.g., CRRA preferences), since we will use a set of equally spaced return states for empirical implementation. Second, we assume that $\phi_{t,T} = \phi_{T-t}$, so that the SDF ratio depends only on the horizon $T - t$. Both assumptions are in effect for the purposes of notational simplification so that we may pool estimates across return-state pairs and expiration dates below.

To derive moment conditions for estimation, we begin by rearranging (39) as

$$\pi_t = \frac{\pi_t^*}{\pi_t^* + \phi_{T-t}(1 - \pi_t^*)}. \tag{57}$$

This equation says how the risk-neutral probability and ϕ_{T-t} together pin down the (unobserved) physical probability, which by definition must be unbiased: $\pi_t = \mathbb{E}_t[\mathbb{1}\{R_{m,0\to T} = \omega_1\} \mid R_{m,0\to T} \in \{\omega_1, \omega_2\}]$. Using this unbiasedness property,

$$\mathbb{E}_t\left[\mathbb{1}\{R_{m,0\to T} = \omega_1\} - \frac{\pi_t^*}{\pi_t^* + \phi_{T-t}(1 - \pi_t^*)} \;\middle|\; R_{m,0\to T} \in \{\omega_1, \omega_2\} \right] = 0. \tag{58}$$

Note that the random variable $\mathbb{1}\{R_{m,0\to T} = \omega_1\}$ is observable as of date T, as it simply indexes whether the terminal index return is equal to ω_1. Thus every value in (58) is in principle observable aside from π_{T-t}, so applying the law of iterated expectations to this equation yields a nonlinear moment condition for ϕ_{T-t} that can be estimated using the generalized method of moments (GMM).

One possible concern with such estimation is the likelihood of price measurement error affecting the measured risk-neutral probabilities in (58) given, for example, market microstructure noise. The GMM framework here, however, allows us to account for this noise without needing to estimate its magnitude separately. To discuss this estimation, we first generalize the notation slightly, and allow for arbitrary return states indexed by j, (ω_j, ω_{j+1}). We then assume that the observed conditional risk-neutral belief $\widehat{\pi}_{t,j}^*$ is measured with additive error with respect to the true value $\pi_{t,j}^*$ used in (58):

$$\widehat{\pi}_{t,j}^* = \pi_{t,j}^* + \epsilon_{t,j}, \tag{59}$$

where $\mathbb{E}[\epsilon_{t+k,j}\, \pi_{t+k',j}^* \mid R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\}] = 0$ for all k, k', and $\epsilon_{t,j}$ follows an MA(q) for some value q.

Using this and then Taylor expanding the observed analogue for the second ter in (58), we obtain

$$\frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} = \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})} + \epsilon_{t,j} + \underbrace{\mathcal{O}\Big(\big(\epsilon_{t,j} + (\phi_{T-t} - 1)\big)^2\Big)}_{\text{higher-order terms}} \tag{60}$$

as $\epsilon_{t,j} \to 0$ and $\phi_{T-t} \to 1$,[19] where the latter limit $\phi_{T-t} = 1$ corresponds to the case of risk-neutrality as seen in (39).

Thus equation (58) can be re-expressed up to higher-order terms as

$$\mathbb{E}_t\left[\mathbb{1}\{R_{m,0\to T} = \omega_1\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} \;\middle|\; R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\} \right] = -\epsilon_{t,j}. \tag{61}$$

The risk-neutral probabilities used on the left side of this equation are now the observable values (inclusive of noise, unlike the ideal values used in (39)). Since $\epsilon_{t,j}$ is assumed to follow an MA(q) as discussed above, we can then form a set of unconditional moments by instrumenting using lagged values of $\widehat{\pi}^*_{t,j}$, for any lags greater than q.

That is, defining the N-dimensional instrument vector

$$Z_{t,j} \equiv \begin{pmatrix} \widehat{\pi}^*_{t-q-1,j} \\ \vdots \\ \widehat{\pi}^*_{t-\overline{q},j} \end{pmatrix} \tag{62}$$

for some $\overline{q} > q$, we can then obtain the time-unconditional orthogonality condition

$$\mathbb{E}\left[\left(\mathbb{1}\{R_{m,0\to T} = \omega_j\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} \mathbb{1}\Big\{R_{m,0\to T} \in \{\omega_j, \omega_{j+1}\}\Big\} \right) Z_{t,j} \right] = 0. \tag{63}$$

This unconditional moment restriction is now amenable to empirical estimation over many expiration dates T, horizons $T - t$, and state pairs j. One can set the sample version of (63) to zero over all pairs $t = \tau_1, T = \tau_2$ such that $\tau_2 - \tau_1 = \kappa$, in order to identify ϕ_κ. One can then stack the moment condition for values of $\kappa = 1, 2, \ldots$, to obtain horizon-dependent risk-price estimates.

For empirical estimation, we define the set of return states (and, by implication, return-state pairs) Ω as discussed in the main text. (A given return state is realized if the excess log return is in a given 2-ppt range.) We use the S&P index options data set (and associated data filters) described in the main text, and we estimate risk-neutral densities as described in the appendix of Lazarus (2022) (see Section 3 for an intuitive discussion of such risk-neutral estimation). We restrict ϕ_{T-t} to be fixed by weeks to expiration, so ϕ_1 is, e.g., the one-week-horizon estimated risk price. Finally, we use the five-day-lagged observed risk-neutral probability $\widehat{\pi}^*_{t-5,j}$ as an instrument in the moment equation for $\widehat{\pi}^*_{t,j}$; following the discussion above, this is equivalent to assuming an MA(4) measurement-noise process and setting $\overline{q} = q + 1 = 5$.

Estimation for the ϕ_κ values shown in Figure 3 is then conducted with two-step GMM. The first-stage weight matrix is $Z'Z/\mathcal{T}$, where Z is the data matrix for the instruments and $\mathcal{T}$ is the number of observations. The second-stage weight matrix is then clustered by blocks of 8 time-adjacent observations. The figure

[19] More formally, one may write the remainder term as $\mathcal{O}((\|\epsilon_{t,j}\| + (\phi_{T-t} - 1))^2)$ as $\|\epsilon_{t,j}\| \to 0$ and $\phi_{T-t} \to 1$, where $\|\epsilon_{t,j}\|$ indexes the bounds on $\epsilon_{t,j}$.

presents the resulting estimates, which are downward-sloping by horizon; see the main text for additional discussion.