The Risk-Return Trade-off Puzzle: Backward- versus Forward-Looking Expected Returns

Stefanos Delikouras*   Matthew Linn†

May 31, 2025

Abstract

The positive relation between risk and expected return is central to financial theory. Empirically, the literature has shown this relationship to be very weak. Using option-based risk-neutral densities and estimated stochastic discount factors, we find that the linear and positive risk-return relation holds in expectation. This result is robust to the choice of discount factor (e.g., non-monotonic, VIX-dependent) used to estimate physical densities. Next, we examine the reason for the negative risk-return relationship in the data. We conclude that the risk-return trade-off holds true in expectation but breaks down empirically due to realized returns being poor proxies of expected returns.

Keywords: expected returns, variance, risk-neutral density, physical density, option-based moments, forward-looking

JEL classification: G10, G13, D84

*corresponding author, Department of Finance, Miami Herbert Business School, University of Miami, email: *sdelikouras@bus.umiami.edu*

†Department of Finance, Isenberg School of Management, University of Massachusetts, email: *mlinn@isenberg.umass.edu*

1 Introduction

The risk-return trade-off is one of the most fundamental concepts in financial economics. Investors' required expected return on investment should increase as a function of an investment's risk. In the case of investments where total risk can be mitigated through diversification, required rates of return should be increasing in *systematic* risk. Idiosyncratic risk can be diversified away and, therefore, does not demand a premium.

Surprisingly, however, empirical evidence of the positive risk-return relationship remains inconclusive. The lack of empirical support is especially puzzling among studies of broad stock market indices that are systematic by their very nature and, therefore, only contain systematic risk. Among these systematic types of assets in very liquid markets, economic theory unequivocally predicts a positive relationship between risk and expected return. In contrast, the empirical evidence from the literature shows a weak (e.g., French et al. (1987); Ghysels et al. (2005); Ludvigson and Ng (2007)) and sometimes negative (e.g., Nelson (1991), Glosten et al. (1993); Whitelaw (1994); Harvey (2001)) risk-return relation.

Estimating the market's risk-return trade-off requires econometric estimates of both the expected return and the variance anticipated by market participants at each point in time. Evidence dating back to Poterba and Summers (1986), Bollerslev et al. (1992), and Ghysels et al. (1996) clearly indicates that market variance is known to be persistent and therefore predictable. Equity market returns, on the other hand, are notoriously difficult to predict either due to econometric issues (e.g., Stambaugh (1999), Ferson et al. (2003), Kostakis et al. (2015)) or due to lack of robustness, replicability, and out-of-sample accuracy (e.g., Ang and Bekaert (2007), Goyal and Welch (2008)). One can argue that econometricians can estimate time-varying anticipated variance of market participants by taking advantage of market variance's predictability. To date, estimates of time-varying expected returns are unreliable, especially at short horizons.

Even with a wealth of econometric tools and high-quality index return data, models based on realized returns do not describe market *expected* returns in real time. Econometricians observe *realized* returns and variances and use them to estimate *expected* returns and variances. However, real-time expected returns of market participants depend on incorporating information in real-time. Time series models of expected returns can tell us ex-post about how past information was incorporated into prices, but there is no guarantee that their predictions accurately describe how investors incorporate information ex-ante.

Motivated by these observations, the purpose of this paper is to provide an understanding of the reason why the risk-return trade-off fails empirically. To this end, we propose new estimates of expected returns and variances for the S&P 500 index derived from option data. Options data are a powerful tool because they can be used to describe entire probability distributions of returns or prices at any point in time. Hence, option-based estimates of expected returns and variances are forward-looking in nature because option data reflect prices of contracts that investors enter into, based upon beliefs about *future* values of the underlying. In this sense, estimated expected returns and variances from options data are superior to models of expected returns using time-series of past returns. Thus, in our options-based methodology, we combine several candidate pricing kernels, e.g., monotonic or non-monotonic, fixed or VIX-dependent, with the risk-neutral density from options data to derive alternative physical measures for market returns. Based on the resulting physical moments, we draw two important conclusions.

First, we show that once we control for VIX-dependence or non-monotonicities in the pricing kernel, the relation between option-based risk (variances) and expected return is linear and positive regardless of the shape of the discount factor (monotonic or non-monotonic, fixed or VIX-dependent). This part of our analysis incorporates elements from the pricing kernel puzzle (see Ait-Sahalia and Lo (2000), Jackwerth (2000)), which finds that nonparametric estimates of the SDF as a function of S&P 500 returns are non-monotonic, contradicting standard models that imply a monotonically decreasing SDF (e.g., CRRA preferences). Under mild assumptions, the physical return density is equal to the risk-neutral density times the inverse of the SDF. Hence, if the SDF is monotonic, higher risk-neutral volatility leads to higher expected returns. However, as suggested in the literature, non-monotonic SDFs can break this mechanical risk-return link.

Second, our analysis highlights that the weak empirical evidence in support of the risk-return trade-off is largely due to the difficulty associated with estimating *expected* returns. Using options-based estimates of return densities and moments, we show that the theoretical risk-return trade-off holds in expectation. Conditional expectations differ substantially from realizations. Thus, trying to estimate conditional expected returns using predictive regressions that are based on backward-looking (lagged) variables cannot fully capture the forward-looking aspect of conditional risk premia.

Specifically, for our empirical analysis, we first extract risk-neutral probability densities from the cross-section of S&P 500 index options prices across four different expirations (1-, 2-, 3-, and 6-month) by calculating the second derivative of option prices with respect to the strike price (Huang and Litzenberger (1989)). This is done based on the methodologies first proposed by Breeden (1979), and subsequently refined by Figlewski (2010) and Linn et al. (2018). This allows for risk-neutral density estimates across different option expirations at each point in time that can vary rapidly to reflect relevant pricing information as it is incorporated into option prices. Nevertheless, the primary distributions of interest for studying risk-return relations are typically *physical* probabilities.

To derive *physical* probabilities, we take advantage of the fact that the stochastic discount factor (SDF) transforms the physical density to the risk-neutral density. Provided that the two densities are absolutely continuous with respect to each other, the inverse of the SDF transforms the risk-neutral density to the physical density. In order to use this important fact, one must take into account several issues regarding the properties of the true stochastic discount factor. To begin with, discount factors are not observable. Secondly, the structural parameters of the discount factor (e.g., intertemporal substitution, risk aversion, discount rate) can be time-varying (e.g., Schreindorfer and Sischert (2022)). Thirdly, contrary to the standard monotonicity assumptions for marginal utility, option-derived empirical pricing kernels can be non-monotonic (e.g., Cuesdeanu and Jackwerth (2018)).

Motivated by these observations, we estimate four alternative specifications for the discount factor: standard power utility (monotonic marginal utility with fixed parameters), power utility with VIX-dependent risk aversion (monotonic marginal utility with time-varying parameters), power utility with quadratic exponent (non-monotonic marginal utility with fixed parameters), power utility with VIX-dependent quadratic exponent (non-monotonic marginal utility with time-varying parameters). These pricing kernels expand the standard power utility specification and allow for non-monotonicities, as well as the possibility of time-varying parameters in the discount factor. Importantly, the effects of the VIX, which for the purposes of this estimation is normalized by its sample average, on risk-aversion are non-linear. Hence, these parametric pricing kernels are quite versatile, and allow for a rich set of risk preferences.

In estimating the pricing kernels, we use a GMM system that combines the moment restrictions in Linn et al. (2018) with rational expectations. Specifically, we estimate the parameters in the four discount factors using the fact that any random variable can be transformed to a standard uniform distribution by evaluating its cumulative distribution function as a function of the random variable. Further, we combine these uniform moments with a rational expectations moment restriction according to which the unconditional average of option-based expected returns should be equal to unconditional averages of realized returns over the sample period. Finally, we multiply the inverse of the four alternative discount factors with the common risk-neutral density to derive alternative option-based physical distributions across the four expirations. These distributions are completely forward-looking since they are based on option prices.

Each of the four discount factors in this study implies a slightly different risk-return relation. For instance, to test the risk-return relation for the baseline monotonic power utility model with fixed risk aversion, we regress the conditional risk-premia from the option-based physical measures on the option-based physical variances. When risk aversion in the monotonic pricing kernel depends on the VIX, the linear risk-return relation includes interactions of option-based physical variances with the VIX. In this case, the risk-return relation is estimated via non-linear least-squares due to the non-linear effects of the VIX.

For the fixed-parameter non-monotonic pricing kernel, the quadratic term implies that the standard risk-return relation is augmented by a third-moment term. Finally, for the VIX-dependent non-monotonic SDF, the risk-return relation must be augmented by both a third-moment term as well as non-linear interactions with the VIX, which are estimated via non-linear least squares. Using the above risk-return relations for each of the four discount factors, we conduct a series of tests with realized and option-based moments. The results of these tests can be summarized as follows.

Consistent with the previous literature, when we regress realized excess returns on realized variances, we find a negative risk-return relation. The relation becomes positive when we replace realized excess returns with fitted excess returns and realized variances with fitted variances. Fitted excess returns are from regressions of realized returns on the price-dividend ratio, the risk-free rate, and the dividend-yield, while fitted variances are from regression of realized variances on past realized variances and the VIX. Despite the positive sign of the risk-return coefficient for fitted returns and variances, this relation is weak as evidenced by the near-zero R^2s since fitted returns are based on backward-looking regressions instead of forward-looking option prices.

When we replace realized returns with option-based conditional risk premia as the dependent variables and realized variances with option-based conditional variances as the explanatory variables in risk-return regressions, the relation is perfectly positive across all discount factors. Hence, the first contribution of the paper is the finding that the strong positive connection between expected returns and variances holds for the option-based physical measures regardless of the functional form of the pricing kernel, e.g., non-monotonicity or VIX-dependence, used to convert the risk-neutral density into a physical one.

Specifically, for any well-motivated, monotonic SDF, the positive relationship between expected returns and expected variances is almost mechanical. However, if risk aversion is time-varying, e.g., dependent on the VIX, or if the SDF is non-monotonic, then the positive risk-return relation is not guaranteed to hold. Nevertheless, we show that even when the SDF is not constrained to be monotonic or fixed-parameter, there is a strong positive relationship between option-based expected returns and variances, once we control for VIX-dependence and non-monotonicities.

Additionally, when the SDF is characterized by non-monotonicities or time-varying parameters it is not clear that the expected return/variance relationship is linear. Using the option-based expected returns and variances, we show that regardless of the shape or VIX-dependence of the pricing kernel, the risk-return trade-off is well approximated by a linear relation once we account for VIX-dependence and non-monotonicities. In addition to these empirical findings regarding the linear and positive risk-return trade-off for option-based moments regardless of non-monotonicities or VIX-dependence, we also provide a theoretical explanation of these results using the framework from actuarial science termed linear and quadratic Esscher transforms (Esscher (1932), Monfort and Pegoraro (2012)).

The finding that non-monotonicities and VIX-dependence in risk aversion do not affect the linear positive relation between risk and return can be explained by the fact that non-mononicities in the discount factor occur over extreme returns, either too low or too high, that are assigned near-zero probabilities. Hence, non-monotonicities of the pricing kernel over zero probability events do not really affect the risk-return trade-off and the resulting physical density. Similarly, even if VIX-dependence of risk aversion affects the entire physical distribution of returns, especially when the pricing kernel is non-monotonic, once we control for VIX-dependence and third moments in risk-return regressions, the positive relation between variances and expected returns is preserved.

The second contribution of the paper is that we use the strong positive linear relation between option-based expected returns and variances as a benchmark, and sequentially replace option-based moments with realized moments in risk-return regressions to identify whether realized returns or realized variances are responsible for the reversal of the positive risk-return regression in the data. In other words, we assume the true risk-return relation is positive as option-based estimates confirm across all discount factors we examine. We then investigate whether realized returns or realized variances are responsible for the reversal of the risk-return relation in the data.

According to this approach, we find that the risk-return relation is preserved when we replace option-based variances with realized or fitted variances. This result is important because it addresses any concerns about a purely mechanical relation between option-based expected returns and variances. Fitted and realized variances are not bound by the risk-return trade-off imposed by the pricing kernels we use to derive the physical density. Hence, the positive relation between forward-looking risk premia and backward-looking fitted or realized variances confirms the theoretical predictions, namely that higher risk implies higher ex-ante expected returns for the stock market. To the contrary, when we replace forward-looking risk premia with backward-looking fitted or realized excess returns as the dependent variable but keep the option-based variances as the explanatory variable, the risk-return relation is much weaker.

We conclude that tests of the risk-return relation fail because realized (or backward-looking fitted) excess returns are vastly different from option-based forward-looking risk premia. Thus, realized returns do not preserve the risk-return relation when used in regressions with option-based variances. In contrast, realized and fitted variances are quite similar to the forward-looking option-based ones. Hence, realized variances are able to maintain the risk-return relation when used in regressions with option-based expected returns. The distinction between realized returns and realized variances in preserving the positive risk-return relation and identifying the exact cause of why the risk-return connection breaks down empirically is the second contribution of this paper.

Overall, this paper adds to the literature that studies the weak empirical evidence for the most important relation in finance, according to which investors should be compensated with higher expected returns for holding riskier assets. To address this discrepancy between theory and empirical data, several alternative approaches have been proposed by the literature. Harvey (2001) employs non-parametric techniques to predict returns, and suggests that the discrepancy in the risk-return relation depends on the specification of conditional variance. Brandt and Kang (2004) model the conditional mean and volatility of stock returns as a latent VAR process. Bali (2008) applies a bivariate GARCH model to a cross-section of industry and size/book-to-market portfolios and finds a positive relation.

Lundblad (2007) uses a long time series (1836-2003) for his tests and shows that accurate testing of the risk-return trade-off requires a lot of data to correctly estimate conditional expectations. Ludvigson and Ng (2007) use dynamic factor analysis to improve the prediction of expected returns, which leads to a positive risk-return correlation. Similarly, Bandi and Perron (2008) find evidence in support of a risk-return trade-off in long-horizon samples, and suggest that short-horizon empirical data affect the analysis of the risk-return relation. Wang et al. (2017) attribute the risk-return discrepancy to volatility not being conditioned on prior gains and losses, a result consistent with reference-dependent preferences.

In relation to the existing literature, our approach is novel in that we highlight the forward-looking aspect of the risk-return relation, which can be tested using option prices. To some extent, this is consistent with the conclusion in previous work (e.g., Lundblad (2007)) that conditional expected returns, which are unobserable, are challenging to estimate. Instead, we show that it is quite straightforward to combine the option-based risk-neutral density with stochastic discount factors to derive forward-looking expected returns under the physical distribution, which can be used to test the risk-return relation.

Our paper is also related to the expanding literature that uses option markets to study investor preferences and assess the accuracy of competing macroeconomic asset pricing models (e.g., Beason and Schreindorfer (2022)). Specifically, Ait-Sahalia and Lo (2000) and Rosenberg and Engle (2000) have motivated a series of works on the monotonicity of discount factors derived from option prices. They document that the option-derived kernels can be increasing over a certain moneyness range. Cuesdeanu and Jackwerth (2018) provide an thorough analysis of the possible causes of the non-monotonic pricing kernel.

Linn et al. (2018) and Kim (2021) introduce conditional estimation of their non-parametric discount factor during periods of high and low VIX values, and, similarly to Barone-Adesi et al. (2020), do not find evidence in support of non-monotonicities in option-derived marginal utility. Nevertheless, this literature also alludes to the potential resolutions of asset pricing puzzles using preference specifications estimated from option prices (e.g., Schreindorfer and Sischert (2022)). We add to this literature by showing that non-monotonicities do not affect the risk-return trade-off unless they are combined with VIX-dependence of risk aversion. This is because non-monotonicities of the pricing kernel occur over extreme events with near-zero probabilities, whereas VIX-dependence affects the entire distribution.

2 Alternative Stochastic Discount Factors

To set the stage for the empirical analysis regarding the risk-return trade-off between realized, backward-looking (from predictive regressions), and forward-looking (option-based) expected returns and variances, we first discuss the various pricing kernels used in deriving the option-based forward-looking expected returns and variances. These discount factors can be broadly classified into monotonic and non-monotonic types. For each group, we assume both constant and time-varying parameters.

2.1 Monotonic Pricing Kernel

The baseline specification used in our tests is power utility defined over stock market wealth W_t:

$$U(W_t) = W_t^{1-\gamma_1}/(1-\gamma_1).$$

Based on the above functional form, the intertemporal marginal rate of substitution between dates t and $t+T$ is given by

$$M_{1,t,t+T}(R_{t,t+T}) = \beta U'(W_{t+T})/U'(W_t) \quad = \quad exp\{log\beta - \gamma_1 ln R_{t,t+T}\}. \tag{1}$$

The constant β is the rate of time preference, and $R_{t,t+T} = \frac{W_{t+T}}{W_t}$ is the return on total equity wealth. The parameter $\gamma_1 > 0$ describes (relative) risk aversion since

$$-\frac{R_{t,t+T}}{M_{1,t,t+T}}\frac{\partial M_{1,t,t+T}}{\partial R_{t,t+T}} = \gamma_1.$$

A natural extension to the standard power-utility pricing kernel of equation (1) is to assume that risk aversion is time-varying depending on observable variables that are known at time t. To this end, Linn et al. (2018) and Kim (2021) introduce conditional estimation of their non-parametric discount factor during periods of high and low VIX values. Given the importance of the VIX in option-pricing and the recent results in Schreindorfer and Sischert (2022), we assume an extension of the monotonic pricing kernel of equation (1) for which γ_1 is time-varying with an explicit dependence on the VIX:

$$M_{2,t,t+T}(R_{t,t+T}) = \beta \frac{U_2'(W_{t+T})}{U_2'(W_t)} = exp\{log\beta - \gamma_1 nvix_{t,t+T}^{\gamma_3} lnR_{t,t+T}\}. \tag{2}$$

Risk aversion in the above discount factor is given by $\gamma_1 nvix_{t,t+T}^{\gamma_3}$, where $nvix_{t,t+T}$ is the VIX ($VIX_{t,t+T}$) normalized by its unconditional average ($\overline{VIX_{t,t+T}}$) and scaled by $\sqrt{T}$, the number of days in 1-, 2-, 3-, and 6-month intervals ($T \approx 30, 60, 90, 180$), divided by $\sqrt{365}$:

$$nvix_{t,t+T} = \frac{VIX_{t,t+T}\sqrt{T}}{\overline{VIX_{t,t+T}}\sqrt{365}}. \tag{3}$$

The $nvix$ specification implies that there is a natural level of risk associated with investing in the stock market, which is captured by the average VIX. When VIX is greater or lower than its average, risk aversion will be affected depending on the signs of γ_1 and γ_3. Similarly, since VIX is a percentage, we opt for the specification $nvix_{t,t+T}^{\gamma_3}$ in equation (2) instead of the term $VIX_{t,t+T}^{\gamma_3}$ to avoid extreme values for risk aversion when VIX is very low (high) and γ_3 is a negative (positive) number. The VIX and the normalized VIX ($nvix$) are known to the investor at time t. Hence, to avoid any look-ahead bias in our empirical analysis, the average VIX, $\overline{VIX}$ in equation (3), is estimated over the 1986-1995 period before the beginning of our sample. The values for the average VIX are 5.356%, 7.818%, 10.186%, and 13.119% for the 1-, 2-, 3-, and 6-month expirations.[1]

Given the positivity of the VIX and $nvix$, a well-defined risk aversion coefficient would require a positive γ_1 parameter in equation (2). The constant γ_3 in equation (2) captures the procyclicality of risk-aversion with respect to the $nvix$. For positive γ_1, if γ_3 is positive (negative), risk aversion is procyclical (counter-cyclical) with respect to the $nvix$. Empirically, we would expect γ_3 to be

[1] The VIX values before 1990 were fitted by estimating the regression of the VIX on the old methodology VIX at the daily frequency. The estimates of this regression are: intercept = 0.240% (t-stat= 1.04), OLS coefficient = 0.987 (t-stat=71.01), R^2 = 96%. Standard errors were calculated with a 30-lag Newey-West correction for autocorellation and heteroscedasticity.

positive so that risk aversion increases with $nvix$. If γ_3 is zero, equation (2) collapses to the utility function with constant risk aversion (equation (1)).

2.2 Non-monotonic Pricing Kernel

One contentious aspect of the option-derived pricing kernel is its monotonicity. Although Linn et al. (2018) do not find evidence in support of non-monotonicities in option-based marginal utility, the existing literature has advocated for U-shaped discount factors (e.g., Ait-Sahalia and Lo (2000), Rosenberg and Engle (2000)). Given the recent findings in Schreindorfer and Sischert (2022) and Driessen et al. (2022), we also allow for non-monotonicities in the pricing kernel. The non-monotonic pricing kernel is described by a power utility function with a quadratic exponent

$$M_{3,t,t+T}(R_{t,t+T}) = exp\{log\beta - \gamma_1 lnR_{t,t+T} - \gamma_2 ln^2 R_{t,t+T}\}. \tag{4}$$

The non-monotonicity of the pricing kernel is captured by the square term $ln^2 R_{t,t+T}$ and the parameter γ_2. A possible interpretation of the non-monotonic model is a power utility specification with a risk-aversion coefficient that depends on the stock market since

$$-\frac{R_{t,t+T}}{M_{3,t,t+T}} \frac{\partial M_{3,t,t+T}}{\partial R_{t,t+T}} = \gamma_1 + 2\gamma_2 lnR_{t,t+T}. \tag{5}$$

Based on the above relation, the coefficient γ_2 determines the procyclicality of the risk aversion coefficient with respect to the returns of the stock market. If γ_2 is positive (negative) then risk aversion is pro-cyclical (counter-cyclical) with respect to the stock market. According to Bakshi and Madan (2007) and Bakshi et al. (2010), another possible interpretation of the quadratic term is that it captures the risk aversion of investors who are shorting the market. Cuesdeanu and Jackwerth (2018) provide a thorough analysis for the possible causes of non-monotonic pricing kernels.

Finally, similar to the VIX-dependent model of equation (2), in addition to non-monotonicities, we introduce a more complicated non-monotonic pricing kernel, which is characterized by dependence of the quadratic preference parameters on the normalized VIX ($nvix$). Specifically, we assume that the discount factor is

$$M_{4,t,t+T}(R_{t,t+T}) = exp\{log\beta - \gamma_1 nvix_{t,t+T}^{\gamma_3} lnR_{t,t+T} - \gamma_2 nvix_{t,t+T}^{\gamma_3} ln^2 R_{t,t+T}\}. \tag{6}$$

The dependence of the above pricing kernel on the normalized VIX ($nvix$) is regulated by the constant γ_3. In this case, risk aversion depends both on the VIX and on stock market returns since

$$-\frac{R_{t,t+T}}{M_{4,t,t+T}}\frac{\partial M_{4,t,t+T}}{\partial R_{t,t+T}} = \gamma_1 nvix_{t,t+T}^{\gamma_3} + 2\gamma_2 nvix_{t,t+T}^{\gamma_3} lnR_{t,t+T}. \tag{7}$$

Because VIX is positive, the coefficient γ_2 in equation (6) determines the procyclicality of risk aversion with respect to the stock market as in the non-monotonic discount factor with fixed parameters of equation (4). If γ_2 is positive (negative), then risk aversion is procyclical (countercyclical) relative to the stock market. Nevertheless, the introduction of a quadratic term complicates the dependence of risk aversion on the normalized VIX, which is given by

$$\gamma_1\gamma_3 nvix_{t,t+T}^{\gamma_3-1} + 2\gamma_2\gamma_3 nvix_{t,t+T}^{\gamma_3-1} lnR_{t,t+T}.$$

In this case, the procyclicality of risk aversion with respect to the VIX also depends on the level of stock market returns.

2.3 Estimation of the Pricing Kernel and the Forward-looking Physical Measure

In estimating the moments under the physical distribution across different pricing kernels, we assume that $Q(R)$ is the cumulative distribution function for gross equity returns, R, under the risk-neutral measure. The risk-neutral density can be estimated from option prices. We also assume that $P(R)$ is the forward-looking continuous cumulative distribution function for equity returns under the physical measure, and $dP(R)/dQ(R)$ is the Radon-Nikodym derivative between the two measures. The forward-looking physical measure is unobservable. However, assuming that the conditions in Linn et al. (2018) hold, the Radon-Nikodym derivative is unique, and is equal to the inverse of the discount factor $M_{t,t+T}(R)^{-1}$

$$dP_{t,t+T}(R)/dQ_{t,t+T}(R) = M_{t,t+T}(R)^{-1}. \tag{8}$$

The forward-looking physical measure can therefore be calculated as the product of the risk-neutral measure with the inverse of the pricing kernel

$$dP_{t,t+T}(R) = M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R). \tag{9}$$

In general, the pricing kernel is also unobservable. However, we can use equations (1), (2), (4), and (6) in a GMM system to estimate the unknown parameters of the discount factor. To this end, we employ two sets of GMM moment conditions. The first one, which is introduced in Linn et al. (2018), is based on the fact that any continuous random variable can be transformed to a standard uniform random variable, $P_{t,t+T}(R) \sim U[0,1]$, by evaluating the cumulative distribution function at the random variable, and thus

$$\mathbb{E}\Big[\Big(\int_0^{R^*_{t,t+T}} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)\Big)^n\Big] = \frac{1}{n+1},\ n = 1, 2, ... \tag{10}$$

The variable $R^*_{t,t+T}$ above is the realized gross equity return between dates t and $t+T$.

The second set of target GMM moments is a rational expectations restriction where the unconditional average of option-based expected returns according to the physical measure are equal to the unconditional averages of realized returns $\big(\mathbb{E}\big[R^*_{t,t+T}\big]\big)$:

$$\mathbb{E}\Big[\int_0^{+\infty} R\ M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)\Big] = \mathbb{E}\big[R^*_{t,t+T}\big]. \tag{11}$$

By equating average realized returns to average expected returns, we establish a link between realized and option-based expected returns. The rational expectations moment condition of equation (11) is a novel component in the GMM estimation of option-based discount factors that has not been used before. Finally, in addition to equation (9), the physical density should integrate to one ($\int_0^{+\infty} dP_{t,t+T}(R) = 1$), and thus the inverse of the pricing kernel should also satisfy

$$1 = \int_0^{+\infty} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R),\ \forall t. \tag{12}$$

Dividing equations (9) and (12) by parts yields a normalized version of the pricing kernel

$$dP_{t,t+T}(R) = \frac{M_{t,t+T}(R)^{-1}}{\int_0^{+\infty} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)} dQ_{t,t+T}(R). \tag{13}$$

Similarly to Bliss and Panigirtzoglou (2004), we estimate the normalized pricing kernel using the above theoretically derived moment conditions to guarantee that the estimated parameters imply well-behaved conditional physical density functions for every date.

After this normalization, the GMM system from equations (10) and (11) becomes

$$\begin{bmatrix} \mathbb{E}\Big[\Big(\frac{\int_0^{R^*_{t,t+T}} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)}\Big)^n\Big] - \frac{1}{n+1}, \; n = 1, 2, ..., \\ \mathbb{E}\Big[\frac{\int_0^{+\infty} R \; M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1} dQ_{t,t+T}(R)}\Big] - \mathbb{E}\big[R^*_{t,t+T}\big] \end{bmatrix}. \tag{14}$$

As in Bliss and Panigirtzoglou (2004), due to this normalization, the discount rate parameter β in equations (1), (2), (4), and (6) cannot be identified and will be dropped. The normalization in equation (14) is equivalent to imposing an additional time-varying parameter in the discount factor as in Schreindorfer and Sischert (2022), which forces every physical density to integrate to one.

We estimate the parameters in the pricing kernels of equations (1), (2), (4), and (6) with an exactly identified single-step GMM, i.e., n is equal to $m - 1$ in equation (14), where m is the number of parameters in the discount factor. For the case that m is one, namely the estimation of the simple power utility model from equation (1), GMM is conducted with an over-identified 2×2 system with one degree of freedom. This over-identification stems from the fact that we force the simple one-parameter discount factor to satisfy both the uniform moments (equation (10)) and the rational expectations (equation (11)) condition.

For all pricing kernels, we use an $m \times m$ diagonal weighting matrix with diagonal elements $\{1, 1, ..., 100\}$. The weighting matrix assigns a weight of one to the uniform moments (equation (10)), and a weight of 100 to the rational expectations condition (equation (11)). This is because the scale of the rational expectations condition is much lower than that of the uniform moment. For the exactly identified GMM systems corresponding to the SDFs of equations (1), (4), (6), the choice of weighting matrix is irrelevant. The weighting matrix mainly affects the estimation of the monotonic fixed-parameter model (equation (1)) whose GMM system is over-identified. The gradient of the GMM function is obtained numerically by differentiating equation (14) with respect to the parameters of each discount factor.

Standard errors, which are calculated according to the methodology in Cochrane (2005), are corrected for autocorrelation and heteroscedasticity using the Newey and West (1987) formula with

12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. For the over-identified GMM system that corresponds to the standard power utility model of equation (1), we can compute the χ^2-test for GMM errors with one degree of freedom. For the exactly-identified GMM system corresponding to the discount factors of equations (2), (4), and (6)), the χ^2-test has zero degrees of freedom.

After estimating the preference parameters via GMM, we can multiply the inverse of the estimated pricing kernel with the option-based risk-neutral density to derive the option-based physical measure. Specifically, the option-based physical density, distribution function, and corresponding physical moments are given by

$$\begin{aligned} dP_{t,t+T}(R)/dR &= \frac{M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}/dR \qquad (15) \\ P_{t,t+T}(R^*_{t,t+T}) &= \frac{\int_0^{R^*_{t,t+T}} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)} \\ \mathbb{E}_t[R_{t,t+T}] &= \frac{\int_0^{+\infty} R\ M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)} \\ var_t(R_{t,t+T}) &= \frac{\int_0^{+\infty}(R-\mathbb{E}_t[R])^2\ M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)} \\ \mathbb{E}_t[(R_{t,t+T}-\mathbb{E}_t[R_{t,t+T}])^3] &= m_{3,t}(R_{t,t+T}) = \frac{\int_0^{+\infty}(R-\mathbb{E}_t[R])^3\ M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}. \end{aligned}$$

The same approach without the term $\frac{M_{t,t+T}(R)^{-1}}{\int_0^{+\infty} M_{t,t+T}(R)^{-1}dQ_{t,t+T}(R)}$ is used to calculate moments of the risk-neutral density. The derivation of the option-based risk-neutral density, which is quite standard in the literature, is discussed in the Online Appendix. Figure OA.1 in the Online Appendix shows the average risk neutral density (RND) obtained from option prices for different expirations. These graphs are very similar to those in Linn et al. (2018).

2.4 Risk-Return Trade-off

Our starting point for deriving the theoretical risk-return relation is the first-order conditions for a representative investor whose marginal utility is given either by the monotonic specifications of equations (1) and (2) or the non-monotonic utility of equations (4) and (6).

2.4.1 Risk-Return Trade-off for the Monotonic Pricing Kernel

For any asset i, the Euler equation for asset returns in excess of the risk-free rate, $R^i_{t,t+T} - R^f_{t,t+T}$, using the pricing kernel of equation (1) reads

$$\mathbb{E}_t\big[\big(R^i_{t,t+T} - R^f_{t,t+T}\big)R^{-\gamma_1}_{t,t+T}\big] = 0.$$

Since $R_{t,t+T}$ is the return on the equity market, which is also a traded asset, the above equation should also hold for $R^i_{t,t+T} = R_{t,t+T}$. Further, using the definition of covariance and a first-order Taylor expansion of the discount factor around $R_{t,t+T} = \mathbb{E}_t[R^{-\gamma_1}_{t,t+T}]^{-\frac{1}{(\gamma_1+1)}}$, we obtain the standard risk-return relation for conditional market risk premia

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \gamma_1 var_t(R_{t,t+T}). \tag{16}$$

Equation (16) states that the conditional risk premium on total equity wealth is proportional to its conditional variance, and establishes a strongly positive relation between risk and return for total equity wealth. Equation (16) is the basis of our empirical estimation, which is conducted via OLS with a restricted zero intercept

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \widetilde{\gamma_1} var_t(R_{t,t+T}) + \epsilon_t. \tag{17}$$

We restrict the intercept to zero, as required by the Euler equations. This approach enables more precise estimates of $\widetilde{\gamma_1}$, undistorted by the inclusion of an intercept that lacks theoretical justification.[2] The most important implication of the reduced-form risk-return relation of equation (17) is that $\widetilde{\gamma_1}$ should be positive and statistically significant. Additionally, the estimated value for $\widetilde{\gamma_1}$ cannot be too large since it is a measure of the risk-aversion parameter γ_1. Finally, a large R^2 in this regression implies that a linear function can accurately capture the risk-return trade-off even if equation (17) is derived from a first-order approximation.

Using the same arguments as above, i.e., definition of covariance and Taylor expansion of the discount factor around $R_{t,t+T} = \mathbb{E}_t[R^{-\gamma_1 nvix^{\gamma_3}_{t,t+T}}_{t,t+T}]^{-\frac{1}{(\gamma_1 nvix^{\gamma_3}_{t,t+T}+1)}}$, the structural risk-reward relation

[2] Including intercepts in our empirical analysis does not qualitatively alter the results.

for the VIX-depended monotonic pricing kernel of equation (2) reads

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \gamma_1 nvix^{\gamma_3}_{t,t+T} var_t(R_{t,t+T}). \tag{18}$$

This is because the normalized VIX, $nvix_{t,t+T}$, is known at time t. The reduced-form counterpart of equation (18) is

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \widetilde{\gamma_1} nvix^{\widetilde{\gamma_3}}_{t,t+T} var_t(R_{t,t+T}) + \epsilon_t, \tag{19}$$

which we estimate with non-linear least squares (NLS). The economic interpretation of equation (19) is very similar to that of equation (17), with the additional feature that the risk aversion parameter is time-varying and depends on the normalized VIX in a non-linear way.

2.4.2 Risk-Return Trade-off for the Non-monotonic Pricing Kernel

The non-monotonic discount factors from equations (4) and (6) imply a risk-return relation different from the one described in equation (16). Specifically, the non-monotonic pricing kernel with VIX-dependence from equation (6) yields the following Euler equation

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = -\frac{covar_t\big(R_{t,t+T}, R_{t,t+T}^{-\gamma_1 nvix^{\gamma_3}_{t,t+T} - \gamma_2 nvix^{\gamma_3}_{t,t+T} lnR_{t,t+T}}\big)}{\mathbb{E}_t[R_{t,t+T}^{-\gamma_1 nvix^{\gamma_3}_{t,t+T} - \gamma_2 nvix^{\gamma_3}_{t,t+T} lnR_{t,t+T}}]}. \tag{20}$$

To facilitate the estimation of equation (20), we take the following steps. First, we approximate the highly non-linear pricing kernel in equation (20) with a second-order polynomial around $R_{t,t+T} = 1$.[3] Hence, equation (20) becomes

$$\begin{aligned}\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \mathbb{E}_t[R_{t,t+T}^{-\gamma_1 nvix^{\gamma_3}_{t,t+T} - \gamma_2 nvix^{\gamma_3}_{t,t+T} lnR_{t,t+T}}]^{-1}\Big(\gamma_1 nvix^{\gamma_3}_{t,t+T} var_t(R_{t,t+T}) \\ +\frac{1}{2}\big(2\gamma_2 nvix^{\gamma_3}_{t,t+T} - \gamma_1^2 nvix^{2\gamma_3}_{t,t+T} - \gamma_1 nvix^{\gamma_3}_{t,t+T}\big) covar_t\big(R_{t,t+T}, (R_{t,t+T}-1)^2\big)\Big).\end{aligned} \tag{21}$$

Second, we approximate the covariance $covar_t\big(R_{t,t+T}, (R_{t,t+T}-1)^2\big)$ with the option-based physical third central moment, $m_{3,t}(R_{t,t+T}) = \mathbb{E}_t\big[R^3_{t,t+T} - \mathbb{E}_t[R_{t,t+T}]^3\big]$. Both this covariance and

[3] For the function $f(y) = y^{-a-blny}$, the following hold: $f(1) = 1$, $f'(1) = -a$, and $f''(1) = -(2b - a^2 - a)$.

the third central moment are driven by the term $\mathbb{E}_t[R^3_{t,t+T}]$.[4] Third, we assume that the term $\mathbb{E}_t[R_{t,t+T}^{-\gamma_1 nvix^{\gamma_3}_{t,t+T} - \gamma_2 nvix^{\gamma_3}_{t,t+T} lnR_{t,t+T}}]$ in equation (21) is approximately one.[5]

Based on the above simplifications, a reduced-form approximation of equation (21), which will be estimated via NLS, is

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \widetilde{\gamma}_1 nvix^{\widetilde{\gamma}_3}_{t,t+T} var_t(R_{t,t+T}) + (\widetilde{\gamma}_2 nvix^{\widetilde{\gamma}_3}_{t,t+T} - 0.5\widetilde{\gamma}_1^2 nvix^{2\widetilde{\gamma}_3}_{t,t+T}) m_{3,t}(R_{t,t+T}) + \epsilon_t, \qquad (22)$$

where we set $\widetilde{\gamma}_2 = (\gamma_2 - 0.5\gamma_1)$. According to the relation between equations (20), (21), and (22), we would expect that $\widetilde{\gamma}_1$ is positive, $\widetilde{\gamma}_2$ is negative, and $\widetilde{\gamma}_3$ is of either sign. Additionally, a large R^2 for this regression implies that a linear specification can accurately explain the risk-return trade-off even for the non-monotonic SDF. Finally, when $\widetilde{\gamma}_3$ is zero in equation (22), we obtain the linearized relation between expected returns, variances, and third moments that corresponds to the fixed-parameter non-monotonic discount factor of equation (4)

$$\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T} = \widetilde{\gamma}_1 var_t(R_{t,t+T}) + \widetilde{\gamma}_2 m_{3,t}(R_{t,t+T}) + \epsilon_t, \qquad (23)$$

where now we set $\widetilde{\gamma}_2 = (\gamma_2 - 0.5\gamma_1 - 0.5\gamma_1^2)$. Equation (23) can be estimated via OLS. We would expect that $\widetilde{\gamma}_1$ is positive, $\widetilde{\gamma}_2$ is negative, and R^2s are large.

Regardless of the functional form of the pricing kernel (monotonic or non-monotonic) or the moments (variances alone or variances and third-moments) in the reduced-form versions of the Euler equations, equations (16), (18), and (20) highlight the basic premise of our tests. Namely, that the risk-return relation is forward-looking and holds in expectation. Hence, if we test this relation using realized returns (and variances) to proxy for the forward-looking and time-varying conditional moments, the results may be spurious.

Importantly, equations (16), (18), and (20) imply that once we control for VIX-dependence and non-monotonicites, the risk-return relation is approximately linear. In our setting, this last

[4] $covar_t(R_{t,t+T}, (R_{t,t+T} - 1)^2)$ is $\mathbb{E}[R^3_{t,t+T} - 2R^2_{t,t+T}] - \mathbb{E}_t[R_{t,t+T}]\mathbb{E}_t[R^2_{t,t+T} - 2R_{t,t+T}]$ and $m_{3,t}(R_{t,t+T})$ is $\mathbb{E}_t[R^3_{t,t+T}] - \mathbb{E}_t[R_{t,t+T}]^3$, where the expectation is with respect to the option-based physical measure.

[5] The sample averages of the term $\mathbb{E}_t[R_{t,t+T}^{-\gamma_1 nvix^{\gamma_3}_{t,t+T} - \gamma_2 nvix^{\gamma_3}_{t,t+T} lnR_{t,t+T}}]$ for the discount factor in equation (6) are 1.25, 1.19, 1.10, and 1.17 for the 1-, 2-, 3-, and 6-month expirations. Its sample variances are 0.078, 0.033, 0.034, and 0.008, respectively. The sample averages of the term $\mathbb{E}_t[R_{t,t+T}^{-\gamma_1 - \gamma_2 lnR_{t,t+T}}]$ for the discount factor in equation (4) are 1.03, 0.99, 0.98, and 0.95 for the 1-, 2-, 3-, and 6-month expirations. Its sample variances are 0.0011, 0.0002, 0.0002, and 0.0006.

implication is the outcome of log-linearizations. Nevertheless, as discussed in Section 6 (equations (28) and (31)), if we assume a log-normal risk-neutral density, the linear relation between risk and return should hold exactly once we control for VIX-dependence and non-monotonicities regardless of the shape of the pricing kernel.

3 Option Data and the Risk-Neutral Density

In this section, we discuss the options data used in our empirical analysis, which examines the risk-return trade-off across realized, predictive regressions-based (backward-looking), and option-based (forward-looking) expected returns and variances.

3.1 Data

For our tests, we use option data with expirations of one, two, three, and six months. Choosing a cross-section of expirations is done for two reasons. First, it helps verify the robustness of our results across expirations. Second, it allows for term structure implications across preference parameters similar to Driessen et al. (2022). We do not include high-frequency options (0DTE's) in our sample so that the time series of consecutive observations in our options dataset spans a long period of time characterized by a wealth of events (Great Recession, Covid, etc.). High-frequency options were introduced after 2015 (2016 for weekly's and 2022 for daily's). Finally, in selecting option contracts, we impose the following filters: non-missing implied volatility, positive volume, and bid price above \$3/8. After imposing these criteria, we also require that outside the $\pm 2\%$ moneyness range there are at least six option contracts for each date, with at least three puts and three calls.

The sample of option contracts on the S&P500 is from OptionMetrics, and is summarized in Table 1. Overall, we obtain 323 RND's for the 1-month expiration, 147 for the 2-month, 83 for the 3-month, and 52 RND's for the 6-month expiration. There are two non-consecutive 2-month observations and eleven non-consecutive 3-month observations at the beginning of the sample for which the risk-neutral distributions could not be estimated. For these dates, there were not enough option contracts that satisfied the sample selection criteria. Calculating autocorrelations for Newey-West standard errors requires consecutive observations. Hence, for the 2- and 3-month options, additional observations had to be deleted. The resulting sample is from January 1996 to December

2022 for 1-month expiration options, May 1998 to November 2022 for 2-month expiration, January 2002 to October 2022 for 3-month expiration, and June 1996 to June 2022 for 6-month options.

To highlight the crucial differences between forward-looking option-based information and backward-looking historical data, we estimate equations equations (17), (19), (22), and (23) using three sets of variables. The first set of variables are forward-looking expected returns and variances from option prices based on equation (15).

For the second set of variables, we follow the existing literature (e.g., Pastor et al. (2008)), and regress realized returns on realized variances to verify whether the fundamental-risk return relation of equation (16) holds in the data. Following Guo and Whitelaw (2006), realized S&P500 variances, $v^2_{t,t+T}$, for each expiration are derived from the following equation

$$v^2_{t,t+T} = \sum_{\tau=t}^{t+\widetilde{T}_t} (r^2_{d,\tau} + 2r_{d,\tau} r_{d,\tau-1}). \tag{24}$$

Above, $r_{d,\tau}$ is the daily S&P500 return. T denotes the 1-, 2-, 3-, and 6-month interval ($T = 30, 60, 90, 180$), and $\widetilde{T}_t$ is equal to the actual number of business days within each 1-, 2-, 3-, and 6-month interval starting at time t ($\widetilde{T}_t \approx 21, 42, 83, 126$).[6] The set of explanatory variables in equations (22) and (23) includes third central moments. Hence, realized third central moments, $s^3_{t,t+T}$, are calculated similar to realized variances

$$s^3_{t,t+T} = \sqrt{\widetilde{T}_t} \sum_{\tau=t}^{t+\widetilde{T}_t} (r^3_{d,\tau} + 3r^2_{d,\tau} r_{d,\tau-1} + 3r_{d,\tau} r^2_{d,\tau-1}), \tag{25}$$

where $r_{d,\tau}$ are daily returns. Following Zhang et al. (2021), we use a factor $\sqrt{\widetilde{T}_t}$ for the third moment such that when we divide equation (25) by $(v^2_{t,t+T})^{2/3}$ from equation (24), the resulting skewness measure is properly scaled.

Finally, for the third set of variables, we use fitted returns and variances. Fitted returns are generated from regressing realized returns on the dividend yield, the lag dividend growth, and the risk-free rate. Similar specifications for fitted returns are suggested in Guo and Whitelaw (2006), Goyal and Welch (2008), and Martin (2017). Fitted variances are from regressing realized returns

[6] The number of business days in each 1-, 2-, 3-, and 6-month intervals varies over time, hence the t-subscript in $\widetilde{T}_t$ of equation (24).

on lagged realized variances and the contemporaneous VIX as in Guo and Whitelaw (2006).

Similarly to fitted variances, fitted third central moments are from regressions of realized third central moments on lagged realized third central moments and the contemporaneous third central moment under the risk-neutral density. We term these fitted moments as backward-looking because they depend on lagged data. In estimating fitted returns, fitted variances, and fitted third moments, we lose one observation due to the use of lag variables. For the 1- and 6-month expirations, when we estimate fitted returns, we lose two observations due to data availability. Regression results for the estimation of fitted moments (returns, variances, third moments) are reported in Table OA.1 of the Online Appendix.

4 Estimation of the Alternative Pricing Kernels

Using the GMM estimation methodology from Section 2 and the RND's from the Online Appendix, we estimate the parameters in the various discount factors from equations (1), (2), (4), and (6). The estimated pricing kernels will help us examine the risk-return relation between forward-looking expected returns and variances under different physical measures. Panel A, Table 2 reports the estimates for the standard monotonic power utility model of equation (1). The risk aversion parameters are positive across expirations, and range from 1.266 (2-month) to 1.524 (6-month). The values of these parameters are consistent with those reported in Bliss and Panigirtzoglou (2004), and imply that risk aversion is stable across expirations.

Panel B, Table 2 reports the results for the monotonic pricing kernel with time-varying risk aversion coefficients that depend on the normalized VIX ($nvix$) from equation (2). γ_1 estimates are positive across all expirations implying positive risk aversion. The estimated value of γ_1 for the 1-month options is 0.518, while the γ_1 estimates for the 2-, 3-, and 6-months options range from 1.295 to 1.661. These discrepancies in γ_1 for the VIX-dependent pricing kernel across expirations can be explained by the coefficient that regulates the dependence of risk aversion on the normalized VIX, γ_3, which is positive for 1-month (1.936), and negative (from -0.817 to -0.559) for the 2-, 3-, and 6-month expirations.

Panel A, Table 3 reports the GMM results for the non-monotonic discount factor of equation (4) with constant parameters. For the 1-month expiration, the linear coefficient γ_1 is positive (0.814) and the quadratic parameter γ_2 is negative (-7.412). This implies that risk-aversion is counter-cyclical with respect to market returns (equation (5)), and that marginal utility is U-shaped. To the contrary, for 2-, 3-, and 6-month expirations, both the linear γ_1 parameter and the quadratic γ_2 coefficients are positive (γ_1 = 1.386-1.672, γ_2 = 0.334-0.779). According to equation (5), this implies that risk aversion for these expirations is pro-cyclical with respect to market returns, and that marginal utility is inverse U-shaped. For the 2-, 3-, and 6-month expirations, the procyclicality of risk aversion with respect to the stock market is counterintuitive. However, this procyclicality is consistent with the counter-cyclicality of risk aversion with respect to the *nvix* for the VIX-dependent model of equation (2) that was shown in Panel B, Table 2.

Finally, Panel B, Table 3 reports estimation results for the non-monotonic pricing kernel with VIX-dependent parameters from equation (6). In this case, the linear coefficient γ_1 is positive (0.129 to 1.991) and the quadratic coefficient γ_2 is negative (-47.470 to -2.143) across all expirations, implying counter-cyclical risk aversion with respect to the stock market and U-shaped marginal utility. The parameter γ_3 that determines the dependence of the linear and quadratic coefficients on the *nvix* is also negative (γ_3 = -9.720 to -4.421). However, negative estimates for the VIX-dependence coefficients γ_3 do not necessarily mean that risk aversion is decreasing in VIX. As shown in equation (7), the overall cyclicality of risk aversion with respect to the *nvix* in the VIX-dependent quadratic discount factor depends on the level of stock market returns.

Figure 1 plots the various pricing kernels based on the GMM results from Table 2 and Table 3. In these graphs, *nvix* in the VIX-dependent discount factors is set equal to its sample average for each expiration, which is a number close to one. With the exception of the 2-, 3-, and 6-month expirations for the fixed-parameters non-monotonic pricing kernel from equation (4), the remaining discount factors appear to be monotonically decreasing for most moneyness values. The fixed-parameters non-monotonic model for 2-, 3-, and 6-month expirations has an inverse U-shape because the estimates for both the linear, γ_1, and the quadratic, γ_2, coefficients are positive. Nevertheless, the increasing part of this pricing kernel corresponds to extremely low values for market returns (e.g., -100% to -50%), which are assigned almost zero probabilities. For realistic market returns (e.g., -50% to 30%), the non-monotonic pricing kernel of equation (4) is decreasing.

Based on the standard errors and t-statistics of the estimates in Table 2 and Table 3, the option-implied pricing kernel is probably monotonic since the quadratic terms are mostly statistically insignificant. Regarding VIX-dependence the evidence most likely points towards VIX-independence with the exception of two instances in Table 3: the 1- and 2-month expirations for the non-monotonic VIX-dependent model of equation (6). Nevertheless, in studying the risk-return trade-off between realized, backward- (fitted), and forward-looking (option-based) expected returns and variances, we use the resulting physical densities and moments from the various estimated pricing kernels, regardless of the statistical significance of their coefficients.

Although many studies of the risk-return trade-off employ conditional tests of the relation between *realized* returns and current or expected variance,[7] our primary tests utilize option-implied estimates of contemporaneous *expected* returns and expected variance, consistent with the Euler equations (16), (18), and (20). Option-based risk-neutral estimates at a given time, t, are constructed using information available at t. We rely on full-sample parameter estimates of the pricing kernel to transform risk-neutral to physical distributions, and to derive fitted values of returns and variances. These full-sample parameter estimates are used for simplicity because parameter estimates vary slowly when calculated with an expanding window. Look-ahead bias is therefore not a concern in our setting.

5 The Risk-Return Trade-off: Empirical Results

In this section, we present the main results regarding the relation between realized and option-based expected returns and variances. The structure of our empirical analysis is the following. We first regress realized returns on realized variances to verify the weak and often negative relation documented in the literature. Next, we regress option-based expected returns on option based variances to verify that the strong positive linear relation between risk and return predicted by the theory holds, regardless of non-monotonicities or VIX-dependencies in the pricing kernel that has been used to convert the risk-neutral density into the physical one. Finally, we sequentially replace option-based moments with realized ones in risk-return regressions to identify whether it is realized returns or realized variances that are responsible for the reversal of the positive risk-return relation.

[7] While the literature is voluminous, some notable examples include Campbell (1987), French et al. (1987), Glosten et al. (1993) and Ludvigson and Ng (2007).

5.1 Risk-Return Regressions: Traditional Approach with Realized and Backward-looking Fitted Returns and Variances

Table 4, Panel A reports summary statistics for realized excess returns, realized variances (equation (24)) and realized third-central moments (equation (25)) across the four option expirations (1-, 2-, 3-, and 6-month). These statistics are sensible and consistent with the literature. Average realized excess returns range from 0.69% to (1-month) to 3.79% (6-month) implying an annualized risk premium of about 8%. Average realized variances range from 0. 26% (1 month) to 1.53% (6 months), resulting in an annualized volatility of 17%.

Average realized third moments are negative across all expirations (-0.14% to -0.01%), which implies negatively skewed distributions. Table 4 also reports summary statistics for the risk-neutral variances of the option-based RND. Average risk-neutral variances range from 0.51% to 2.67% which translates into an annualized risk-neutral volatility of about 24%. Finally, the average normalized VIX (*nvix*) is a number around one (1 to 1.42), since the average VIX during the period 1986-1995 is used as normalization. Yet, the *nvix* exhibits large deviations especially for short-term expirations (1- and 2-months).

Table 4, Panel B reports correlations between realized excess returns, variances, skewness, the risk-neutral variance, and the *nvix* across option expirations. As expected, the correlation between the risk-neutral variances and the *nvix* is about 0.9 (0.84 to 0.93). Most importantly, the results in Panel B of Table 4 offer a glimpse at the risk-return trade-off puzzle in that the correlation between realized excess returns and realized variances are negative and economically significant with a monotonic pattern across all expirations: -0.19 (1-month) to -0.63 (6-month). This finding confirms the puzzling negative relation between realized risk and return documented by the existing literature. Panel B of Table 4 also offers a potential resolution of this negative relation puzzle once we substitute realized variances with forward-looking ones. Specifically, the correlation of realized excess returns with option-based risk-neutral variances is positive, albeit small (0.03 to 0.21).

Table 5 reports results for the risk-return relation based on realized excess returns and variances by estimating equation (17) without an intercept as prescribed by the structural Euler equation for risk premia (equation (16)). A strong and economically meaningful risk-return relation would imply that the $\widetilde{\gamma}_1$ coefficient in equation (17) is positive and statistically significant and that the R^2

of the regression would be large. Contrary to the theoretical predictions of the Euler equation (16), the relation between realized excess returns and realized variances in Panel A of Table 5 is negative across all expirations ($\widetilde{\gamma}_1$: -2.128 to -0.818). This negative relation between realized returns has a reasonable fit since the R^2s are non-zero (2.08% to 20%).

Interpretingly, when we replace realized excess returns with backward-looking fitted returns or realized variances with backward-looking fitted variables in regression (17), then the relation between realized returns and fitted variances (Panel B, Table 5), fitted excess returns and realized variances (Panel C, Table 5) or fitted excess returns and fitted variances (Panel D, Table 5) becomes positive. In these regressions, backward-looking fitted returns are based on regressing realized returns on the price-dividend ratio, the risk-free rate, and dividend growth. Similarly, backward-looking fitted variances are based on regressions of realized variances on past realized variances and the VIX.[8] This findings are the first pieces of evidence that the positive risk-return relation holds in expectation but not necessarily in realization. Nevertheless, across all these regressions, the R^2s are very low (-15% to 19%) implying a weak risk-return relation contrary to the predictions of the Euler equation (16).[9]

5.2 Risk-Return Regressions: Forward-looking Expected Returns from the Monotonic Discount Factors

The results in Table 5 indicate that the risk-return trade-off of equation (16), which dominates the financial economics paradigm, should hold in expectation, and that when realized returns and realized variances are used to test this relation the results are economically meaningless (negative relation between risk and returns, low R^2s). To the contrary, when we substitute realized returns with fitted returns from a regression with the price-dividend ratio, the risk-free rate, and dividend growth, or when we swap realized variances with fitted variances from a regression with past realized variances and the VIX, the risk-return regressions align with the theoretical predictions (positive relation between risk and returns).

[8]The results from the backwards-looking regressions as well as summary statistics for the backward-looking expected returns, variances, and third moments are reported in Table OA.1 of the Online Appendix.

[9]Restricting the intercept to zero in our regressions can result in negative R-squared. This occurs because imposing this restriction causes the regression estimate to deviate from the orthogonal linear projection that minimizes the sum of squared errors. In cases where the restricted estimate provides a particularly poor fit, the sum of squared errors may exceed the total sum of squares, leading to a negative R-squared.

For the next set of tests, we estimate equation (17) using purely forward-looking expected returns and variances from the option-based physical density that corresponds to the monotonic discount factor from equation (1). Similarly, we test equation (19) with forward-looking expected returns and variances from the option-based physical density that corresponds to the monotonic discount factor from equation (2) with VIX-dependent parameters.

Table 6 reports summary statistics for the option-based expected returns and variances for the fixed-parameter (Panel A) and VIX-dependent monotonic discount factors (Panel B) across the different expirations. Average risk premia for the two pricing kernels (0.67% to 3.80%) are identical to average realized returns from Table 4. This is because the GMM conditions that estimate the two pricing kernels set average option-based expected returns equal to average realized returns as highlighted in equation (11). Although variances are not target moments in the GMM system that estimates the various pricing kernels, the averages of option-based variances are strikingly consistent (0.42% to 1.99%) between fixed-parameter and VIX-dependent monotonic discount factors across all expirations.

The most important finding in Table 6 is the almost perfect positive correlation between forward-looking risk premia and variances across all maturities for both the fixed-parameter (0.98-0.99) and the VIX-dependent (0.85-0.96) monotonic discount factors. These results show that the positive risk-return relation for the monotonic discount factor is not affected by time variation in risk aversion induced by the VIX. Importantly, the near perfect correlations highlight the strong positive relation between option-based returns and variances, which is the starting point of our analysis.

5.2.1 Fixed-parameter Monotonic Discount Factor

Table 7, Panel A reports results from regressing forward-looking risk premia on forward-looking variances from the fixed-parameter discount factor of equation (1). Consistent with the correlation results in Table 6, the regression estimates of the risk aversion coefficients $\widetilde{\gamma}_1$ are positive (1.641-1.941) and highly statistically significant (t-stats: 23.73-50.40). Further, the magnitude of these estimates are consistent with the risk aversion estimates from Panel A of Table 2 (γ_1: 1.266-1.524). Importantly, the R^2 of these regressions is almost one (96%-98%). The results in Panel A confirm that for the monotonic discount factor of equation (1), there is a perfect linear relation between risk premia and variances as suggested by theory.

The results in Panel A serve as a benchmark for the risk-return regression tests. In a world where the risk-return relation holds perfectly in expectation, to what extent could tests of realized or backward-looking fitted excess returns and variances identify this positive relation? To answer this question, we sequentially replace option-based moments with realized moments in risk-return regressions to identify whether it is realized returns or realized variances that are responsible for the reversal of the positive risk-return relation in the data. The results are reported in Panels B through E of Table 7.

The strong linear and positive relation between risk and returns is preserved in Panel B of Table 7 where we regress forward-looking risk premia on backward-looking fitted variances. The risk aversion estimates ($\widetilde{\gamma}_1$: 1.837-2.637) are higher than those from regressing forward-looking option-based risk premia on forward-looking option-based variances in Panel A, yet the former estimates are statistically significant (t-stats: 10.98-20.99). The large R^2s (60%-92%) indicate a strong linear relation.

The results are similar for the risk aversion estimates ($\widetilde{\gamma}_1$: 1.302-1.678; t-stats: 4.16-7.79) when we regress forward-looking risk-premia on realized variances in Panel C of Table 7, although the R^2 decreases (-48% to 38%) especially for the 6-month expiration. Finally, when we regress backward-looking fitted (Panel D) or realized excess returns (Panel E) on forward-looking option-based variances, the risk-return relation remains positive and statistically significant although R^2s decrease substantially (-0.61 to 5.91%) relative to Panel A, where we use forward-looking option-based risk premia as the dependent variable.

Overall, the results from Table 7 confirm the strong positive relation between option-based forward-looking risk premia and variances implied by the fixed-parameter monotonic discount factor of equation (1). In addition, the very high R^2 suggests a strong linear relationship. This relation holds even if we replace the option-based variances with backward-looking fitted or realized variances. Importantly, when we replace forward-looking risk premia with backward-looking fitted or realized excess returns in the regression of equation (17), the results are much weaker than when we replace forward-looking variances with realized or backward-looking fitted in the same regression.

These findings suggest that tests of the risk-return relation using realized returns and variances fail mostly because realized (or backward-looking fitted) excess returns are vastly different from option-based forward-looking risk premia. To the contrary, realized and backward-looking fitted variances are quite similar to the option-based ones, and preserve the positive relation in risk-return regressions with option-based expected returns.

The results in Table 7 are illustrated in Figure 2 that plots option-based forward-looking risk premia, backward-looking fitted returns, and realized returns against forward-looking, backward-looking, and realized variances across the 1-, 2-, 3-, and 6-month expirations for the discount factor of equation (1). When the option-based forward-looking risk premia are used as a dependent variable to test the risk-return relation, then the relation is linear, positive, and statistically significant regardless of whether the explanatory variable is the forward-looking, backward-looking or realized variance. Across all horizons, the plots when using forward-looking estimates of variance and expected returns are remarkably linear. Furthermore, the R^2 in Panel A of Table 7 ranges from 98.49% at the 1-month horizon to 96.68% at the 3-month horizon. This suggests that a linear specification is sufficient for analyzing the risk-return trade-off.

To the contrary, when realized or backward-looking fitted excess returns are used as the dependent variable in the risk-return regression of equation (17), the relation is quite weak regardless of the explanatory variable (forward-looking or backward-looking variance), and even turns negative when realized excess returns are regressed on realized variances. This pattern, which is consistent across all option expirations, shows that realized returns cause the reversal of the positive risk-return relation because they are poor proxies of conditional expected returns.

5.2.2 VIX-dependent Monotonic Discount Factor

Table 8 reports results for the forward-looking risk and return trade-off when using the VIX-dependent monotonic discount factor of equation (2). The relation between expected returns and variances for this pricing kernel is estimated via non-linear least squares based on equation (19).

Panel A in Table 8 reports the results when forward-looking option-based risk premia are regressed on option-based variances multiplied by the VIX normalized by its 1986-1995 average ($nvix$). In these regressions, risk aversion parameters, $\widetilde{\gamma}_1$, are positive (0.767-2.101) and statistically significant (t-stats: 11.87-26.19). The VIX-dependence parameters, $\widetilde{\gamma}_3$, are positive (1.703; t-stats: 18.96) for 1-month expiration and negative (-0.823 to -0.468 t-stats: -25.65 to -13.63) for the rest of the expirations and also statistically significant. These estimates imply pro-cyclical, with respect to the VIX, risk-aversion for 1-month expiration and counter-cyclical, with respect to the VIX, risk aversion, for the rest of the expirations. These estimates are consistent with the GMM estimates from Panel B, Table 2.

Importantly, the R^2s in Panel A,Table 8 are large (89%-98%), which implies a strong risk-return relation even if the risk aversion parameter in the monotonic pricing kernels is time varying. This is a novel finding that indicates that VIX-dependence in risk aversion does not alter the linear and positive risk-return trade-off implied by monotonic discount factors regardless of the sign of VIX-dependence (pro- or counter-cyclical).

However, similarly to the methodology in Table 7, the regressions in Panel A, Table 8 are not the main test of the risk-return trade-off. Instead, these results serve as a benchmark of risk-return regressions. In a world where, conditional on the VIX, the risk-return relation holds perfectly in expectation, to what extent could tests of realized or backward-looking fitted returns and variances identify this positive relation?

To this end, Panels B and C in Table 8 replace forward-looking variances with backward-looking fitted and realized variances, while Panels D and E replace forward-looking risk premia with backward-looking fitted and realized excess returns. Finally, Panels F and G replace both forward-looking risk premia and variances with backward-looking fitted and realized moments. These tests also help identify whether it is realized returns or realized variances that cause the breakdown of the positive risk-return relation.

The results from these tests are very similar to those from Table 7 for the fixed-parameter monotonic pricing kernel. Specifically, when we replace forward-looking variances with fitted variances, the risk-return relation of equation (19) remains positive and statistically significant ($\widetilde{\gamma}_1$: 0.909-2.575; t-stats: 5.05-18.91), VIX dependence is also statistically significant, albeit positive for the 1-month expiration ($\widetilde{\gamma}_3$: 1.788; t-stats: 13.69) and negative of the rest ($\widetilde{\gamma}_3$: -1.001 to -0.172; t-stats: -8.07 to -1.02), and the R^2s are large (48%-98%).

When we replace forward-looking variances with realized ones from equation (24) in Panel C of Table 8, the strong risk-return relation for the VIX-dependent monotonic pricing kernel from equation (2) is economically meaningful (R^2: 89%) only for the 1-month expiration ($\widetilde{\gamma}_1$: 0.384; t-stats: 3.16). For the remaining expirations although $\widetilde{\gamma}_1$ is positive (1.250-1.477; t-stats: 3.24-6.67), the R^2s are negative (-439% to -72%) implying a weak risk-return relation.

The results are similar in Panels D through F when we replace forward-looking expected risk premia with fitted and realized excess returns. Although risk aversion estimates are positive (0.998-3.036), the R^2s are low (-2.10% to 11.76%). Finally, similar to the results from Table 5, in Panel G of Table 8 when we regress realized excess returns on realized variances interacted with the VIX, the resulting $\widetilde{\gamma}_1$ coefficients are negative (-3.986 to -1.176; t-stats: -2.59 to -1.27) and the R^2s are low (-7% to 6%).

Similarly to the results in Table 7, Table 8 shows that even if we allow the discount factor to exhibit time-varying risk aversion that depends on the VIX, we cannot use realized returns to test the risk-return relation of equation (19), which is a forward-looking relation. Further, although backward-looking fitted variances, and to a lesser extent realized variances, are a good measure of option-based forward-looking variances, realized and backward-looking fitted returns are very poor proxies of forward-looking expected returns. Thus, using realized or fitted returns to test the risk-return relation delivers weak results.

5.3 Risk-Return Regressions: Forward-looking Expected Returns from the Non-monotonic Discount Factors

For the next set of tests, we examine the more complex risk-return relation of the non-monotonic discount factors, which is described in equations (22) and (23). To set the stage for these tests, Table 9 reports the summary statistics for risk premia, variances, and third moments implied by the non-monotonic discount factor with fixed parameters from equation (4) in Panel A. Panel B reports summary statistics for the moments of the non-monotonic pricing kernel with VIX-dependence from equation (6). Both pricing kernels yield almost identical average risk premia (0.69% to 3.79%), which are also equal to average realized excess returns from Table 4 and average risk premia from monotonic models from Table 6. This because the GMM estimation aligns average option-based expected returns to average realized returns (equation (11)).

With the exception of the 1-month maturity, the average variances from the VIX-dependent non-monotonic pricing kernel (0.34% - 2.23%) are lower than the average variances from the fixed-parameter non-monotonic discount factor (0.36% - 1.32%). This is because the quadratic parameter estimates, γ_2, in the fixed-parameter pricing kernel for the 2-, 3-, and 6-month expiration are positive (Table 3, Panel A) implying an inverse U-shaped marginal utility (Figure 1). To the contrary, the quadratic parameter estimates for the VIX-dependent pricing kernel are negative.

Despite these discrepancies, forward-looking physical variances are generally similar across the two discount factors, even though this moment is not part of the GMM estimation. This result, which is also consistent with the variances from the monotonic discount factors in Table 6, confirms that although variances are not part of the GMM, all discount factors generate similar average variances.

Compared to the previous results, the novel moment in Table 9 is the third central moment, which is essential for the risk-return trade-off of the non-monotonic discount factors (equations (22) and (23)). The third central moment is negative across the two non-monotonic discount factors (-0.01 to -0.32) implying negatively skewed physical distribution. The most important result in Table 9 are the correlations between the first three moments for the non-monotonic pricing kernels. Specifically, consistent with the results for the monotonic discount factors from Table 6, for the fixed-parameter non-monotonic pricing kernel of equation (4), the correlations between the option-based risk premia and variances are strongly positive (0.60-0.98), while the correlations between option-based risk premia and third central moments are negative (-0.82 to -0.28).

To the contrary, for the VIX-dependent non-monotonic discount factor of equation (6), the correlations between the option-based risk premia and variances are strongly negative (-0.66 to -0.38), whereas the correlations between option-based risk premia and third central moments are positive (0.31 to 0.52). These findings show that non-linearities in the discount factor combined with VIX-dependence can generate completely different implications for physical moments than non-linearities alone. Interestingly, for both the fixed-parameter and the VIX-dependent non-monotonic pricing kernels, option-based variances are negatively correlated to option-based third central moments (-0.93 to -0.44).

5.3.1 Fixed-parameter Non-monotonic Discount Factor

Table 10 reports the linear regression results of equation (23) for the fixed-parameter non-monotonic model of equation (4). In Panel A, we report regressions of option-based risk premia on option-based variances and third central moments. Consistent with the GMM estimates in Panel A, Table 3, the relation between risk-premia and variances is positive ($\widetilde{\gamma}_1$: 1.576-2.029) and significant (t-stat: 10.42-48.22) across all expirations, while the relation between risk-premia and third moments is negative for the 1-month expiration ($\widetilde{\gamma}_2$: -4.379; t-stat: -1.18) and positive for all other expirations ($\widetilde{\gamma}_2$: 0.844-3.418; t-stats: 4.48-24.65).

Importantly, the fit of equation (23) is accurate across all expirations (R^2: 85%-99%). This is a novel finding, as we show that non-monotonicities in the discount factor do not affect the positive linear risk-return relation in expectation. This is because non-monotonicities occur over extreme values for index returns (either to low or too high) that are assigned an almost zero measure (Figure 1). Nevertheless, similarly to the tests for the monotonic discount factor, regressing option-based risk premia on option-based variances and third moments is not our main test. These regressions serve as a benchmark to identify whether such a relation can be uncovered using realized moments.

To this end, in Panel B, Table 10 we substitute option-based forward-looking variances and third moments with backward-looking fitted variances and third-moments from the regressions in Table OA.1 of the Online Appendix. The results are very similar to those in Panel A with the option-based moments. Specifically, the estimates for the linear coefficient $\widetilde{\gamma}_1$ are positive and statistically significant (1.158-3.029; t-stats: 5.61-32.75). Further, estimates for the quadratic parameter $\widetilde{\gamma}_2$ are mostly negative (-9.580 to 5.727; t-stats: -3.66 to 4.56). Nevertheless, there is a sign discrepancy between the GMM estimates for γ_2 in Panel A, Table 3 (positive for 2- and 3-month expiration) and the corresponding estimates in Panel B, Table 10 (negative for 2- and 3-month expiration). However, the fit of regression (23) with backward-looking fitted variances and third moments is impressive (R^2: 84%-89%).

In Panel C of Table 10, we replace option-based variances and third-moments with realized ones from equations (24) and (25), but maintain the option-based risk premia as the dependent variables. In this case, the fit of equation (23) deteriorates but remains reasonable for 1-, 2-, and 3- month expiration (R^2: 27%-53%), while the variance parameter $\widetilde{\gamma}_2$ is positive (1.540-2.029) and statistically significant (6.44-14.20) across all expirations.

In sum, the results in Panels A through C of Table 10 highlight the fact that the risk-return-third moment relation from equation (23) holds for option-based risk premia, variances, and third moments. This relation remains strong when we replace option-based variances and third moments with backward-looking fitted moments from Table OA.1 of the Online Appendix. The risk-return relation deteriorates, but with reasonable R^2s for most expirations, when we replace option-based variances and third moments with realized ones.

To the contrary, when we replace option-based risk premia with realized or backward-looking fitted excess returns as the dependent variable in Panels D through G, the fit of the linear risk-return relation of equation (23) deteriorates significantly (R^2: -2% to 14%). Nevertheless, regression coefficients for realized and backward-fitted variances are positive across most cases ($\widetilde{\gamma}_1$: -0.735 to 4.268; t-stats: -1.48 to 4.39) including those in Panel G, where we regress realized returns on realized variances and skewness.

Hence, consistent with the results from the monotonic pricing kernels, the main implication of Table 10 is that realized returns cannot be used to identify the positive risk-return relation even if this relation is implied by the non-monotonic discount factors (equation (23)). This is because the risk-return relation holds in expectation and realized or backward-looking fitted excess returns are quite different from the forward-looking option-based risk premia. To the contrary, backward-looking fitted variances and to a lesser extent realized variances are similar to forward-looking variances from option prices, and thus, they preserve the linear and positive risk-return relation.

5.3.2 VIX-Dependent Non-monotonic Discount Factor

We obtain similar results in Table 11 where we estimate the more complex relation between risk and returns of equation (22) for the VIX-dependent non-monotonic discount factor of equation (6). When we regress option-based risk premia on option-based variances and third-central moments in Panel A, consistent with the GMM estimates from Panel B ofTable 3, the variance coefficients are positive ($\widetilde{\gamma}_1$: 0.049 to 3.903; t-stats: 1.32-13.92), and the third moment and *nvix* parameters are negative ($\widetilde{\gamma}_2$: -7.315 to -182.010, $\widetilde{\gamma}_3$: -3.462 to -9.776).

The R^2s of the risk-return regression for the VIX-dependent non-monotonic discount factor are large (69%-79%) with the exception of the 6-month expiration (29%). This finding suggests that once we control for VIX-dependence (*nvix*) and non-monotonicities (third moment), the relation between risk and returns can be accurately described by a linear relation even if the pricing kernel is non-monotonic and the parameters are time-varying.

Results are similar in Table 11, Panel B when we replace option-based variances and third moments with backward-looking fitted ones. However, when we regress option-based risk premia on realized variances and third moments (Panel C) or when we replace option-based risk premia with backward-looking fitted or realized returns (Panels D to F), the strength of the non-linear risk-return-third moment relation breaks down (R^2: -336.15% to 19.76%), although the variance coefficient is mostly positive ($\widetilde{\gamma}_1$: -0.395 to 5.090).

Contrary to the results in Panel G of Table 7 and Table 8 for the monotonic discount factors and Panel G of Table 10 for the fixed-parameter non-monotonic pricing kernel, when we regress realized excess returns on realized variances and third-moments based on the non-linear risk-return trade-off with VIX-dependence from equation (22) in Panel G of Table 11, the fit of the model is quite good (R^2; 29%-48%) and the variance parameter is positive ($\widetilde{\gamma}_1$: 0.171 to 1.078) albeit mostly statistically insignificant (t-stat:0.92-3.96). Although the signs of the third-moment and *nvix* parameters are not consistent with the GMM estimates from Panel B in Table 3, Panel G of Table 11 shows that VIX-dependence strengthens the risk-return relation for realized moments.

Overall, the tests in Tables Table 6 through Table 11 identify a positive relation of forward-looking option-based risk premia with backward-looking fitted and to a lesser extend with realized variances. To the contrary, when we replace forward-looking risk premia with backward-looking fitted or realized excess returns as the dependent variable but keep the option-based variances as the explanatory variable, the risk-return relation is much weaker than when we replace forward-looking variances with realized or backward-looking fitted ones in the same regression and keep option-based risk premia as the dependent variable.

We conclude that tests of the risk-return relation fail because realized (or backward-looking fitted) excess returns are vastly different from option-based forward-looking risk premia. Thus, realized returns do not preserve the risk-return relation when used in regressions with option-based variances. In contrast, realized and fitted variances are quite similar to the forward-looking option-based ones. Hence, realized variances are able to maintain the risk-return relation when used in regressions with option-based expected returns. This distinction between realized returns and realized variances in preserving the positive risk-return relation and identifying the exact cause of why the risk-return connection breaks down empirically lies at the core of our exercise.

5.4 Linearity of the Risk-Return Trade-off

In this section, we conduct additional tests to shed light on the risk-return relation between option-based risk premia and higher moments for the different discount factors. Specifically, we show that once we account for VIX-dependence in risk aversion, the risk-expected return relation is linear and positive across all four pricing kernels considered in this study.

5.4.1 Standard Risk-Return Trade-off for Non-standard Discount Factors

As shown in Section 2, each discount factor used in our study implies a different risk-return relation. The baseline power utility model of equation (1) yields the standard positive risk-return relation of equation (17). Introducing VIX-dependent risk aversion in the standard monotonic discount factor as in equation (2) generates VIX-dependence in the risk-return trade-off as shown in equation (19). Finally, augmenting the standard monotonic power utility model to include quadratic terms either with (equation (6)) or without (equation (4)) VIX-dependence introduces third moments in the risk-return trade-off (equations (22) and (23)). In this section, we examine whether the standard linear risk-return expression of equation (17) can capture the relation between option-based risk premia and variances from the non-standard discount factors of equations (2), (4), and (6).

Surprisingly, according to Panel A in Table 12, the standard linear risk-return relation of equation (17) does an almost perfect job in describing the relation between option-based risk premia and variances from the fixed-parameters non-monotonic discount factor of equation (4). The risk aversion estimates for $\widetilde{\gamma}_1$ (0.748-2.203; t-stats: 2.19-65.78) are positive, and the R^2s for the 1-, 3-, and 6-month expirations are quite strong, ranging from 84% to 98%. To the contrary, the R^2 for the 2-month expiration is relatively low (22%).

According to Panel B, Table 12, the standard risk-return relation in equation (17) cannot accurately describe the relation between option-based risk-premia and option-based variances for the VIX-dependent monotonic model of equation (2) with R^2s ranging from 24% to 84%, although the risk aversion estimates are positive and significant ($\widetilde{\gamma}_1$: 1.214-3.551; t-stats: 3.36-13.68).

The results deteriorate significantly in Panel C, where we use the standard risk-return relation of equation (17) to describe the relation between risk-premia and variances from the VIX-dependent non-monotonic discount factor of equation (6). Although the risk aversion estimates are positive (0.228-1.268; t-stats: 1.96-3.27), the R^2s are negative (-342% to -39%). Nevertheless, when instead of equation (17), we use the VIX-dependent monotonic relation of equation (19) to describe the relation between risk premia and variances for the non-monotonic VIX-dependent discount factor in Panel D of Table 12, the fit of the model improves significantly (R^2: 27%-67%).

The results from these tests are graphically summarized by Figure 3, which plots option-based expected returns and option-based variances across expirations and pricing kernels. Consistent with the regression results from Table 12, the graphs in column (c) of Figure 3 show that the linear risk-return relation of equation (17) implied by the standard fixed-parameter monotonic discount factor of equation (1) can capture the relation between risk premia and variances for the non-monotonic discount factor of equation (4) quite well. This is because non-monotonicities in the discount factor occur over values for returns that are assigned near-zero probabilities, either in the far-left or far-right of the distribution of returns as shown in Figure 1. Thus, the standard-risk return relation is minimally affected by these non-monotonicities.

To the contrary, as shown by the graphs in columns (b) and (d) of Figure 3, the standard risk-return relation of equation (17) does a worse job of capturing the relation between risk premia and variances for discount factors with VIX-dependent risk aversion, especially for non-monotonic ones (equation (6)). This is because VIX-dependence affects the entire distribution of returns and hence, distorts the standard risk-return trade-off, especially if VIX-dependence is also combined with non-monotonicities in the pricing kernel.

These results suggest that non-monotonicities in the pricing kernel alone cannot affect the standard linear risk-return relation unless they are combined with VIX-dependence, which tends to amplify the effects of these non-monotonicities. Nevertheless, once we account for VIX-dependence by interacting psychical variances with the VIX as in equations (19) and (22), then the relation between option-based expected returns and variances can be well approximated by a linear function regardless of the shape of the pricing kernel as highlighted by the large R^2's in Panel A of Table 8, Table 10, and Table 11 and the graphs in column (e) of Figure 3.

6 Theoretical Explanation of Results

To explain the empirical findings in this paper, we can use some standard results from the normal distribution. Most of the theoretical results in this section are for log-returns and normal distributions. These results can be applied to gross returns using either the approximation $logR \approx R - 1$ or the exact formulas from the log-normal distribution.

In actuarial science, the baseline power utility of equations (1) and (2) with the normalization of equation (13) is called a linear Esscher transform (Esscher (1932)). More recently, Monfort and Pegoraro (2012) introduced the concept of second-order Esscher transform, which is similar to the non-monotonic pricing kernels from equations (4) and (6). Within this framework, as shown in the Online Appendix, if the risk-neutral density for log-returns is normal with mean $\widetilde{r}^f_{t,t+T}$ and variance $var^{rnd}_t(lnR_{t,t+T})$, the monotonic discount factors from equations (1) and (2) would result in a physical density for log-returns which is normal with mean and variance given by

$$\mathbb{E}_t[lnR_{t,t+T}] = \widetilde{r}^f_{t,t+T} + \gamma_{1,t} var^{rnd}_t(lnR_{t,t+T}) \tag{26}$$

$$var_t\big(lnR_{t,t+T}\big) = var^{rnd}_t(lnR_{t,t+T}), \tag{27}$$

where $\gamma_{1,t}$ is risk aversion that can be time-varying. Hence, when the risk-neutral density is normal, linear discount factors shift expected returns by the product of the risk aversion with the variance, but leave the risk-neutral variance unaltered.

From equations (26) and (27), we conclude that

$$\mathbb{E}_t[lnR_{t,t+T}] = \widetilde{r}^f_{t,t+T} + \gamma_{1,t} var_t\big(lnR_{t,t+T}\big) \tag{28}$$

Equation (28) explains the perfect results for the risk-return relation implied by the monotonic pricing kernels in Panel A of Table 7 and Table 8. These results confirm that once we properly account for the possible time variation of risk aversion (Table 8), there is a strong linear and positive relation between variances and expected returns for monotonic discount factors regardless of VIX-dependence in risk aversion.

Contrary to the results for the linear pricing kernel (equations (26) and (27)), when the risk-neutral density is normal, the quadratic pricing kernel from equations (4) and (6) alters both the mean and the variance of log-returns. Specifically, as shown in the Online Appendix, if the risk-neutral density for log-returns is normal, the quadratic discount factor yields a normal physical distribution for log-returns with mean and variance

$$\mathbb{E}_t[lnR_{t,t+T}] = \frac{\widetilde{r}^f_{t,t+T} + \gamma_{1,t} var^{rnd}_t(lnR_{t,t+T})}{1 - 2\gamma_{2,t} var^{rnd}_t(lnR_{t,t+T})} \tag{29}$$

$$var_t\big(lnR_{t,t+T}\big) = \frac{var^{rnd}_t(lnR_{t,t+T})}{1 - 2\gamma_{2,t} var^{rnd}_t(lnR_{t,t+T})}. \tag{30}$$

In this case, the physical variance is less than the risk-neutral one only if the quadratic discount factor has a negative quadratic parameter $\gamma_{2,t}$.

For the fixed-parameter non-monotonic pricing kernel of equation (4), where $\gamma_{1,t}$ and $\gamma_{2,t}$ are constant, the denominators above are approximately equal to one because $var_t^{rnd}(R_{t,t+T})$ is a relatively small number (0.51% to 2.67% from Table 4), especially for short expirations. Hence, equations (29) and (30) can be approximated by equations (26) and (27) quite accurately. This explains the results in Table 12 where the moments from fixed-parameter non-monotonic pricing kernel perfectly satisfy the linear risk-return relation of equation (16).

To the contrary, for the non-monotonic discount factor with VIX-dependent parameters of equation (6) the denominator in equations (29) and (30) is not necessarily equal to one due to VIX-dependence. In this case, we obtain

$$\mathbb{E}_t[lnR_{t,t+T}] = \frac{\widetilde{r}^{f}_{t,t+T}}{1 - 2\gamma_2 nvix^{\gamma_3}_{t,t+T} var_t^{rnd}(lnR_{t,t+T})} + \gamma_1 nvix^{\gamma_3}_{t,t+T} var_t(lnR_{t,t+T}). \tag{31}$$

The above equation shows that the quadratic pricing kernel with VIX-dependent parameters implies a highly non-linear relation between risk and expected returns because the physical variance ($var_t(lnR_{t,t+T})$) is also multiplied by the $nvix$, which, in turn, is highly correlated with risk-neutral variances ($var_t^{rnd}(lnR_{t,t+T})$). Equation (31), combined with the fact that the empirical risk-neutral density is not actually normal, explains why the standard linear risk-return relation of equation (16) does not apply to the non-monotonic VIX-dependent discount factor of equation (6) as shown in Panel C, Table 12. Instead, for this pricing kernel, we would require to account for both VIX-dependent coefficients and high-order moments to capture the linear and positive risk-return relation as evidenced by equation (21) and the findings in Panel A, Table 11.

Regarding the relation between risk-neutral and physical variances, equation (27) can explain their perfect comovement for the case of monotonic pricing kernels (equations (1) and (2)) as evidenced in Table 6. Table 9 also documents that there is a strong relation between risk-neutral and physical variances even for the non-monotonic pricing kernels (equations (4) and (6)). This is because the denominator in equation (30) is a number around one implying that the relation between risk-neutral and physical variance for non-monotonic pricing kernels is similar, albeit less perfect, to that predicted by equation (27) for monotonic models.

In sum, as shown in equations (26), (27), (29) and (30), when the RND is normal, the linear pricing kernel only affects expected returns but not the risk-neutral variance, whereas the quadratic discount factor affects both expected returns and variance. However, the quadratic effect on the risk-neutral variance is small in magnitude, since $\gamma_{2,t}$ in the denominator of equation (30) is multiplied by $var_t^{rnd}(lnR_{t,t+T})$, which is a small number. Thus, the denominator in equation (30) is very close to one. This finding explains why, as documented in Table 6 and Table 9, the physical variances for the various pricing kernels are highly correlated to the risk-neutral one.

On the other hand, expected returns vary both across pricing kernels and from the risk-neutral density due to the $\gamma_{1,t}var_t^{rnd}(lnR_{t,t+T})$ term in equation (26) and in the numerator of equation (29). Specifically, for the non-monotonic VIX-dependent pricing kernel of equation (6), the non-linear relation of equation (29) is amplified by VIX-dependence. Hence, contrary to the rest of the discount factors for which option-based physical expected returns and variances are positively correlated, expected returns from the non-monotonic and VIX-dependent pricing kernel are negatively correlated to option-based variances as shown in Table 9. However, as shown in Panel A of Table 11, once we properly account for VIX-dependence and non-monotonicities according to equation (20), the relation between risk premia and variances becomes linear and positive even for the VIX-dependent non-monotonic model .

7 Conclusion

We empirically study the risk-return trade-off for the aggregate stock market. In testing this relation, the existing literature usually regresses realized returns on realized variances, and finds a negative relation between risk and returns. This is contrary to the core of the most basic finance theory that investors should be compensated for bearing risk. Instead of using realized returns and variances to empirically investigate the risk return trade-off, we use option prices, the risk-neutral density, and a set of alternative discount factors to derive option-based physical distributions for stock market returns. To generate the physical distribution, we employ four power-utility-based pricing kernels, whose exponents are either linear (monotonic) or quadratic (non-monotonic), and whose risk aversion parameters are either fixed or dependent on the VIX.

Contrary to the findings of the existing literature that uses realized returns and variances, when we regress option-based risk premia on option-based variances from the different pricing kernel, the relation between risk and returns is positive and strong even after controlling for the VIX or non-monotonicities. The core of our empirical analysis uses the strong positive link between option-based expected returns and variances as a benchmark, and sequentially replaces option-based moments with realized ones in risk-return regressions to identify whether it is realized returns or realized variances that are responsible for the reversal of the positive risk-return regression in the data.

We find that the positive risk-return trade-off remains robust even after swapping option-based variances with fitted variances from regressions of realized variances on the VIX and on lagged variances, and to a lesser extent with realized variances. Conversely, we show that the risk-return trade-off severely deteriorates and turns negative when we replace option-based risk premia with realized or fitted returns using the risk-free rate, dividend growth, and the price-divided ratio. Our results suggest that the risk-return trade-off is purely forward-looking and, thus, we cannot test it using realized or backward-looking fitted returns as a proxy for risk premia because these two sets of variables are different. Our main conclusion is that the empirical study of the risk-return trade-off requires purely forward-looking moments such as those derived from option prices.

Another contribution of our analysis is that non-monotonicities or VIX-dependence of the discount factor alone cannot affect the linear relation between option-based risk premia and variances. Instead, when non-monotonicities are combined with VIX-dependence in the pricing kernel, the standard risk-return relation between risk premia and variances requires additional non-linear terms that depend on the VIX and on third moments. This is because non-monotonicities in the pricing kernel occur over extreme values of returns that are assigned almost zero probabilities. To the contrary, VIX-dependence of the discount factor affects the entire physical distribution. Even in this case, however, once we properly account for non-monotonicities and VIX-dependence in risk-return regressions, the forward-looking option-based moments imply a strong linear risk-return trade-off that satisfies the most fundamental theoretical relation in finance. Namely, that higher risk, in expectation, compensates investors with higher returns.

References

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

Alexiou, L., A. Goyal, A. Kostakis, and L. Rompolis (Forthcoming). Pricing event risk: Evidence from concave implied volatility curves. *Review of Finance*.

Ang, A. and G. Bekaert (2007). Stock return predictability: Is it there? *Review of Financial Studies 20*, 651–707.

Bakshi, G. and D. B. Madan (2007). Investor heterogeneity, aggregation, and the non-monotonicity of the aggregate marginal rate of substitution in the price of market-equity. *Working paper*.

Bakshi, G., D. B. Madan, and G. Panayotov (2010). Returns of claims on the upside and the viability of u-shaped pricing kernels. *Journal of Financial Economics 97*, 130–154.

Bali, T. G. (2008). The intertemporal relation between expected returns and risk. *Journal of Financial Economics 87*(1), 101–131.

Bandi, F. M. and B. Perron (2008). Long-run risk-return trade-offs. *Journal of Econometrics 143*(2), 349–374.

Barone-Adesi, G., N. Fusari, A. Mira, and C. Sala (2020). Option market trading activity and the estimation of the pricing kernel: A bayesian approach. *Journal of Econometrics 216*, 430–449.

Beason, T. and D. Schreindorfer (2022). Dissecting the equity premium. *Journal of Polictical Economy 130*, 2203–2222.

Birru, J. and S. Figlewski (2012). Anatomy of a meltdown: The risk neutral density for the S&P 500 in the fall of 2008. *Journal of Financial Markets 15*(2), 151–180.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), 637–654.

Bliss, R. R. and N. Panigirtzoglou (2004). Option-implied risk aversion estimates. *Journal of Finance 59*(1), 407–444.

Bollerslev, T., R. Y. Chou, and K. F. Kroner (1992). Arch modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics 52*(1-2), 5–59.

Brandt, M. W. and Q. Kang (2004). On the relationship between the conditional mean and volatility of stock returns: A latent var approach. *Journal of Financial Economics 75*(2), 217–257.

Breeden, D. and R. Litzenberger (1978). Prices of state-contigent claims implicit in option prices. *Journal of Business 51*(4), 621–51.

Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics 7*(3), 265–96.

Campbell, J. Y. (1987). Stock returns and the term structure. *Journal of financial economics 18*(2), 373–399.

Cochrane, J. H. (2005). *Asset pricing*. Princeton, NJ: Princeton University Press.

Cuesdeanu, H. and J. C. Jackwerth (2018). The pricing kernel puzzle: Survey and outlook. *Annals Finance 14*, 289–329.

Driessen, J., J. Koëter, and O. Wilms (2022). Horizon effects in the pricing kernel: How investors price short-term versus long-term risks. *Jounral of Financial and Quantitative Analysis*.

Esscher, F. (1932). Onn the probability function in the collective theory of risk. *Skandinavisk Aktuarietidskrift 15*(3), 175–195.

Ferson, W. E., S. Sarkissian, , and T. T. Simin (2003). Spurious regressions in financial economics? *The Journal of Finance 58*(4), 1393–1413.

Figlewski, S. (2010). Estimating the implied risk-neutral density for the us market portfolio. *Volatility and Time Series Econometrics: Essay in Honor of Robert F. Engle*, 323–353.

French, K. R., G. Schwert, and R. F. Stambaugh (1987). Expected stock returns and volatility. *Journal of Financial Economics 19*(1), 3–29.

Ghysels, E., A. C. Harvey, and E. Renault (1996). Stochastic volatility. *Handbook of Statistics 14*, 119–191.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2005). There is a risk-return trade-off after all. *Journal of Financial Economics 76*(3), 509–548.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance 48*(5), 1779–1801.

Goyal, A. and I. Welch (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Guo, H. and R. F. Whitelaw (2006). Uncovering the risk–return relation in the stock market. *Journal of Finance 61*(3), 1433–1463.

Harvey, C. R. (2001). The specification of conditional expectations. *Journal of Empirical Finance 8*(5), 573–637.

Huang, C.-F. and R. H. Litzenberger (1989). *Foundations for Financial Economics*. New York, NY: North-Holland.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *The Review of Financial Studies 13*(2), 433–451.

Kim, H. J. (2021). Characterizing the conditional pricing kernel: A new approach. *Working Paper*.

Kostakis, A., T. Magdalinos, and M. P. Stamatogiannis (2015). Robust econometric inference for stock return predictability. *Review of Financial Studies 28*, 1506–1553.

Linn, M., S. Shive, and T. Shumway (2018). Pricing kernel monotonicity and conditional infromation. *Review of Financial Studies 31*(2), 493–531.

Ludvigson, S. C. and S. Ng (2007). The empirical risk–return relation: A factor analysis approach. *Journal of Financial Economics 83*(1), 171–222.

Lundblad, C. (2007). The risk return tradeoff in the long run: 1836–2003. *Journal of Financial Economics 85*(1), 123–150.

Martin, I. (2017). What is the expected return on the market. *Quarterly Journal of Economics 132*(2), 367–433.

Monfort, A. and F. Pegoraro (2012). Asset pricing with second-order esscher transforms. *Journal of Banking & Finance 36*(6), 1678–1687.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–370.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*, 703–8.

Pastor, L., M. Sihna, and B. Swaminathan (2008). Estimating the intertemporal risk–return tradeoff using the implied cost of capital. *Journal of Finance 63*(6), 2859–2897.

Poterba, J. M. and L. H. Summers (1986). The persistence of volatility and stock market fluctuations. *The American Economic Review 76*(5), 1142–1151.

Rosenberg, J. V. and R. F. Engle (2000). Empirical pricing kernels. *Journal of Financial Economics 64*, 341–72.

Schreindorfer, D. and T. Sischert (2022). Volatility and the pricing kernel. Working Paper.

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics 54*, 375–421.

Wang, H., J. Yan, and J. Yu (2017). Reference-dependent preferences and the risk–return trade-off. *Journal of Financial Economics 123*(2), 395–414.

Whitelaw, R. F. (1994). Time variations and covariations in the expectation and volatility of stock market returns. *Journal of Financial Economics 49*(2), 515–541.

Zhang, Z., M. He, Y. Zhang, and Y. Wang (2021). Realized skewness and the short-term predictabuluty for aggregate stock market volatility. *Economic Modelling 103*, 1–13.

Figures

Figure 1 Pricing Kernels

This figure shows the monotonic and non-monotonic pricing kernels of equations (1), (2), (4), and (6) across the 1- (Panel A), 2- (Panel B), 3- (Panel C), and 6-month (Panel D) expirations. Estimation of the pricing kernels is reported in Table 2 and Table 3. In these graphs, we set the value of the normalized VIX ($nvix$) equal to its sample average for pricing kernels with VIX-dependent parameters (Panels B and D). The sample is from January 1996 to December 2022 for 1-month expiration options, May 1998 to November 2022 for 2-month expiration, January 2002 to October 2022 for 3-month expiration, and June 1996 to June 2022 for 6-month expiration options.

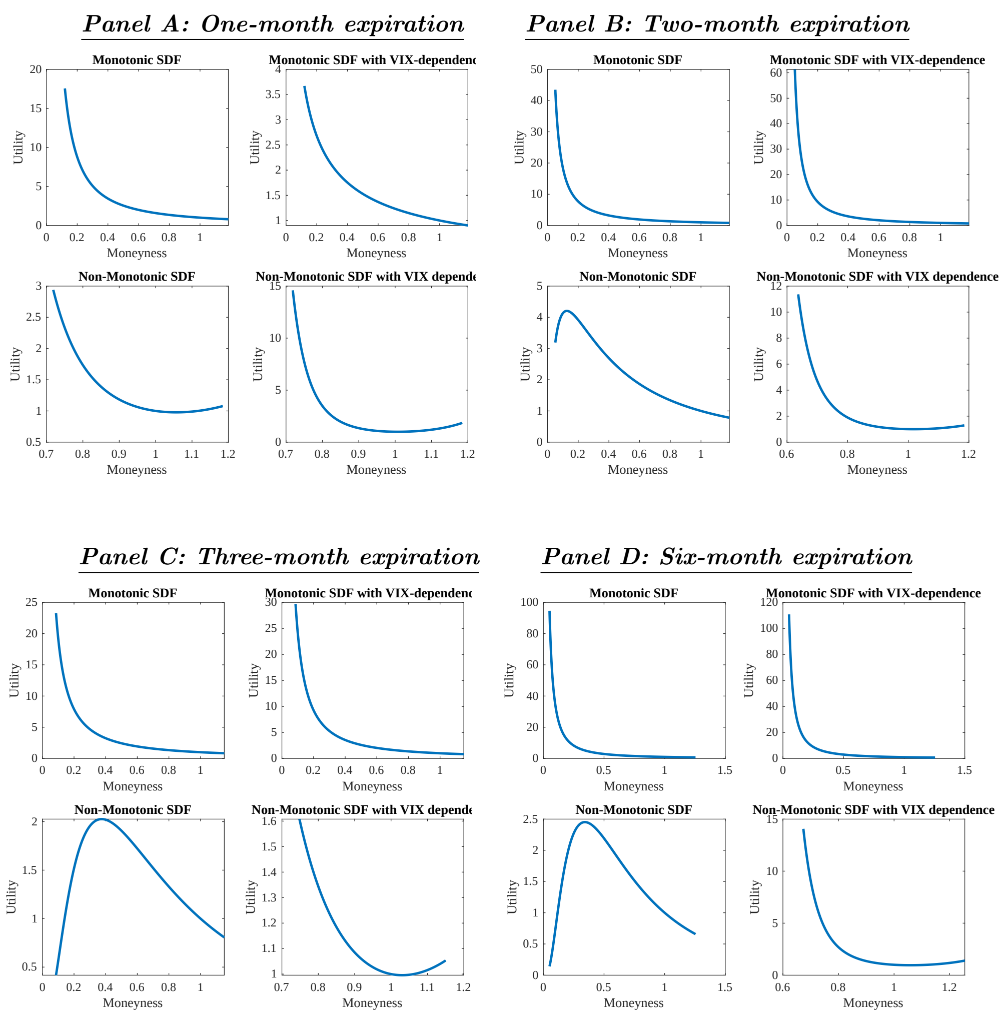

Figure 2 Scatter Plots of Expected Returns and Variances: Linear Pricing Kernel

This figure shows scatter plots of variances (x-axis) and returns (y-axis). Panels A, B, C, and D report results for the 1-, 2-, 3-, and 6-month expirations, respectively. In each Panel, graph (a) is the scatter plot of forward-looking risk premia, $\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking variances, $var_t(R_{t,t+T})$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the estimated pricing kernel of equation (1). Forward-looking S&P500 variances are also calculated from the option-based physical density (equation (15)) using the risk-neutral density and the estimated pricing kernel of equation (1). The estimation of the pricing kernel is reported in Table 2. Graph (b) is the scatter plot of forward-looking risk premia on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared. These regressions are reported in Table OA.1 of the Online Appendix. Graph (c) is the scatter plot of forward-looking risk premia on realized variances, $v^2_{t,t+T}$. Realized S&P500 variances are calculated according to equation (24). Graph (d) is the scatter plot of backward-looking fitted excess returns, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - R^f_{t,t+T}$, on forward-looking variances. Backward-looking fitted S&P500 returns, $\widehat{r}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate. These regressions are reported in Table OA.1 of the Online Appendix. Graph (e) is the scatter plot of backward-looking fitted excess returns on backward-looking fitted variances. Graph (f) is the scatter plot of backward-looking fitted excess returns on realized variances. Graph (g) is the scatter plot of realized returns on forward-looking variances. Graph (h) is the scatter plot of realized returns on realized variances. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$.

Panel A: One-month expiration

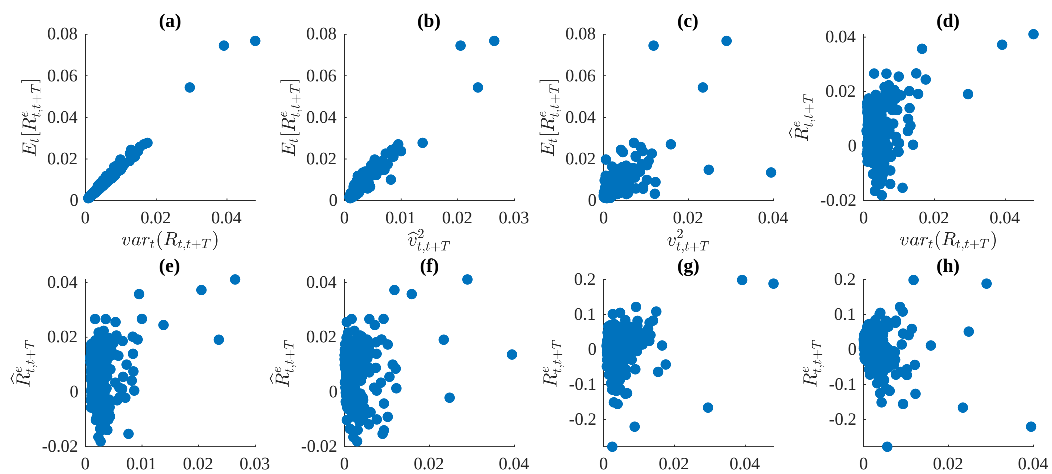

Panel B: Two-month expiration

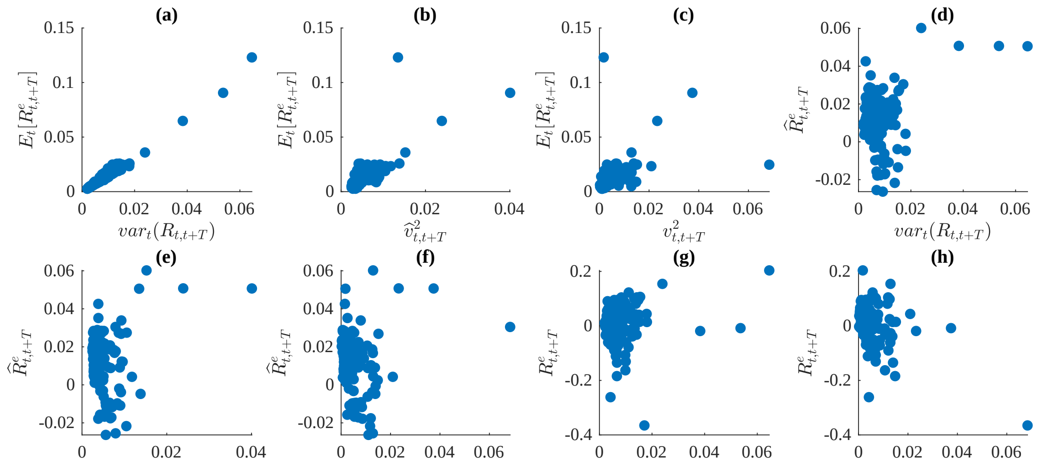

Panel C: Three-month expiration

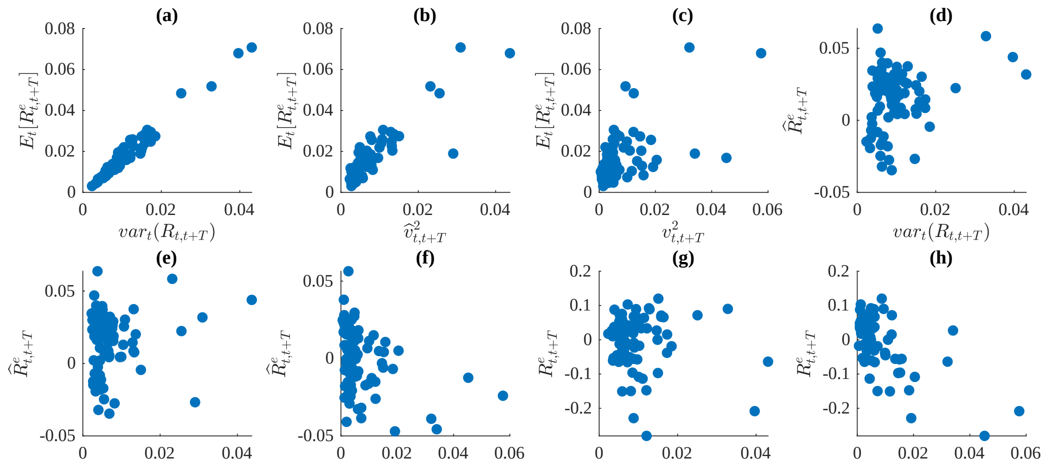

Panel D: Six-month expiration

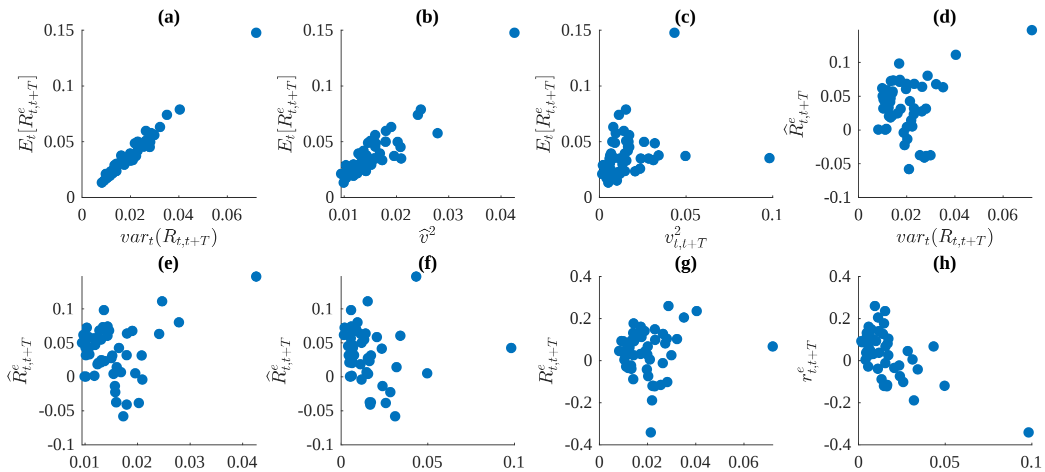

Figure 3 Scatter Plots of Expected Returns and Variances: Option Based Moments

This figure shows scatter plots of variances (x-axis) and returns (y-axis). Panels A, B, C, and D report results for the 1-, 2-, 3-, and 6-month expirations, respectively. All graphs show the scatter plot of forward-looking risk premia, $\mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking variances, $var_t(R_{t,t+T})$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the estimated pricing kernels. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. Forward-looking S&P500 variances are also calculated from the option-based physical density. The estimation of the pricing kernels is reported in Table 2 and Table 3. In the figure, pricing kernels are labeled according to their exponents ($\gamma_1 r$, $\gamma_1 nvix^{\gamma_3} r$, $\gamma_1 r + \gamma_2 r^2$, $\gamma_1 nvix^{\gamma_3} r + \gamma_2 nvix^{\gamma_3} r^2$). Graph (a) shows moments for the fixed-parameter linear pricing kernel of equation (1). Graph (b) depicts moments for the VIX-dependent linear pricing kernel of equation (2). Graph (c) reports moments for the fixed-parameter quadratic pricing kernel of equation (4). Graph (d) shows moments for the VIX-dependent quadratic pricing kernel of equation (6). Graph (e) shows moments for the VIX-dependent quadratic pricing kernel with physical variances multiplied by $nvix^{\widetilde{\gamma}_3}$, where the parameter $\widetilde{\gamma}_3$ is from Table 12, Panel D.

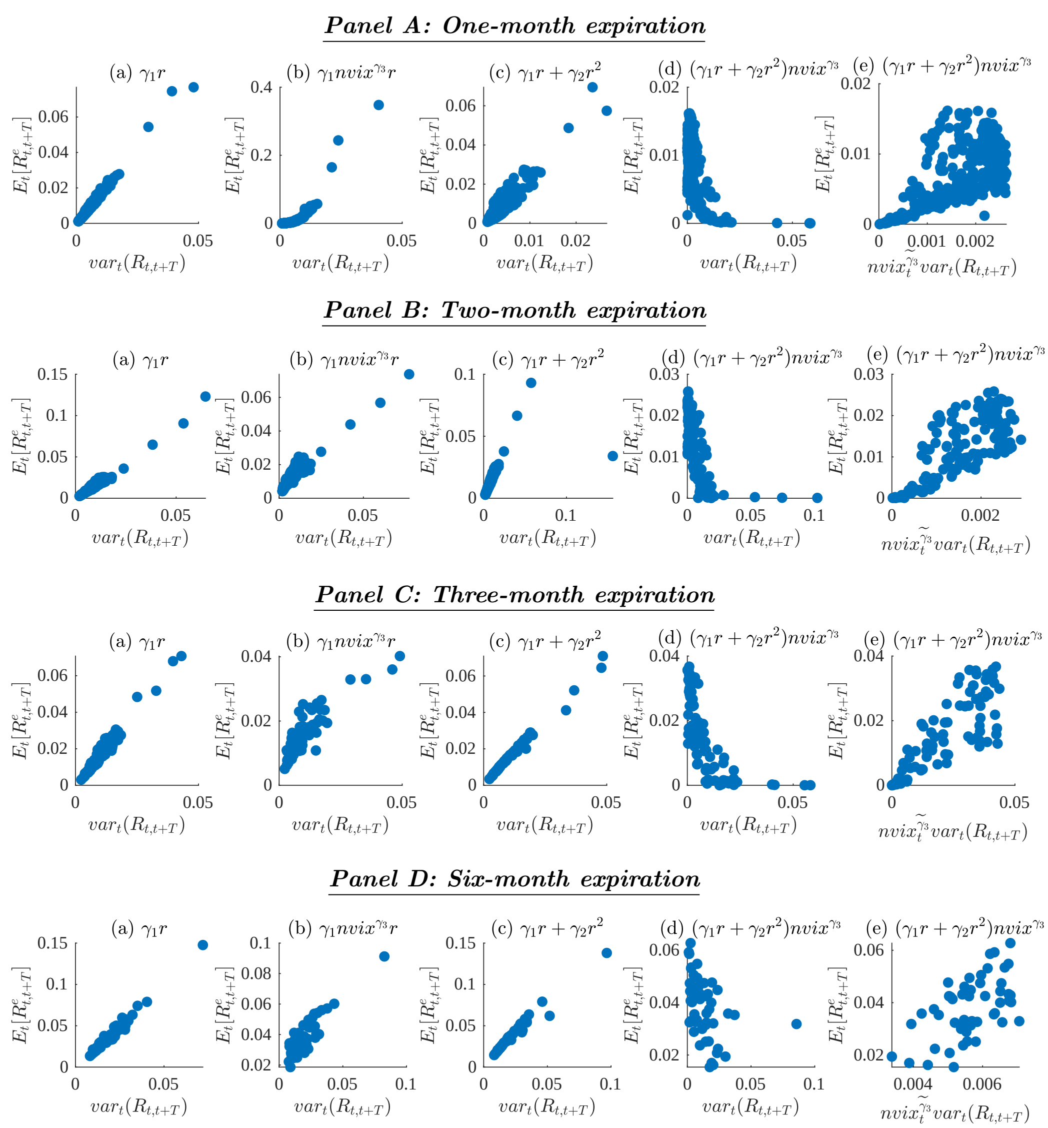

Tables

Table 1 Sample of Option Contracts

This table reports summary information for the S&P500 option contracts from OptionMetrics used to derive the forward-looking density for S&P500 under the physical measure. In selecting option contracts, we impose the following filters: non-missing implied volatility, positive volume, and bid price above \$3/8. We also require that there are at least six option contracts for each date, with at least three puts and at least three calls outside the $\pm 2\%$ moneyness range. There are two non-consecutive 2-month observations and eleven non-consecutive 3-month observations at the beginning of the sample for which the risk-neutral density could not be estimated. For these dates, there were not enough option contracts that satisfied the sample selection criteria. Further, a sample of consecutive observations is required to calculate autocorrelations for Newey-West standard errors in the GMM estimation. Hence, for the 2- and 3-month contracts additional observations had to be deleted from the beginning of the sample. The resulting sample is from January 1996 to December 2022 for 1-month expiration, May 1998 to November 2022 for 2-month, January 2002 to October 2022 for 3-month, and June 1996 to June 2022 for 6-month expiration.

	1-month	2-month	3-month	6-month
Num. of Calls	38,577	14,329	6,141	3,368
Num. of Puts	46,420	18,980	8,432	3,574
Total	84,997	33,309	14,573	6,942
Num. of Expiration Dates (without filters)	323	161	107	52
Num. of Expiration Dates (with filters)	323	159	96	52
Num. of Expiration Dates (with filters and consecutive non-missing observations)	323	147	83	52

Table 2 GMM Estimation of the Option-based Monotonic Pricing Kernel

This table reports GMM results for the fixed-parameter (Panel A) and the VIX-dependent (Panel B) option-based monotonic pricing kernels of equations (1) and (2). γ_1 is the risk aversion parameter and γ_3 is the VIX-dependence coefficient. For the estimation, VIX-dependence is captured by the normalized VIX ($nvix$), which is the VIX divided by its 1986-1995 average (to avoid look-ahead bias) and appropriately scaled for each expiration. The GMM moment conditions are given in equation (14). t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. χ^2, dof, and p are the χ^2 test, degrees of freedom, and p-value that all target moments are jointly zero. GMM is the minimized value of the GMM objective. The sample is from January 1996 to December 2022 for 1-month expiration options, May 1998 to November 2022 for 2-month expiration, January 2002 to October 2022 for 3-month expiration, and June 1996 to June 2022 for 6-month expiration options.

Panel A: Fixed-parameter monotonic discount factor

	1-month	2-month	3-month	6-month
γ_1	1.349	1.266	1.284	1.524
	(4.77)	(4.65)	(4.43)	(6.83)
χ^2	0.21	0.01	0.01	0.01
dof	1	1	1	1
p	0.64	0.90	0.91	0.89
GMM	$3.99e^{-05}$	$6.01e^{-06}$	$9.08e^{-06}$	$2.49e^{-05}$

Panel A: VIX-dependent monotonic discount factor

	1-month	2-month	3-month	6-month
γ_1	0.518	1.439	1.295	1.661
	(0.31)	(1.03)	(5.12)	(2.13)
γ_3	1.936	-0.558	-0.817	-0.813
	(0.47)	(-0.12)	(-0.11)	(-0.15)
χ^2	-	-	-	-
dof	0	0	0	0
p	-	-	-	-
GMM	$1.45e^{-13}$	$2.64e^{-13}$	$3.42e^{-17}$	$1.73e^{-15}$

Table 3 GMM Estimation of the Option-based Non-monotonic Pricing Kernel

This table reports GMM results of the fixed-parameter (Panel A) and VIX-dependent (Panel B) option-based non-monotonic pricing kernels of equations (4) and (6). γ_1 is the parameter for the linear term and γ_2 is the parameter of the quadratic term in the discount factor. γ_3 is the VIX-dependence coefficient. For the estimation, VIX-dependence is captured by the normalized VIX ($nvix$), which is the VIX divided by its 1986-1995 average (to avoid look-ahead bias) and appropriately scaled for each expiration. The GMM moment conditions are given in equation (14). t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. χ^2, dof, and p are the χ^2 test, degrees of freedom, and p-value that all target moments are jointly zero. GMM is the minimized value of the GMM objective. The sample is from January 1996 to December 2022 for 1-month expiration options, May 1998 to November 2022 for 2-month expiration, January 2002 to October 2022 for 3-month expiration, and June 1996 to June 2022 for 6-month expiration options.

Panel A: Fixed-parameter non-monotonic discount factor

	1-month	2-month	3-month	6-month
γ_1	0.814	1.386	1.439	1.672
	(0.57)	(1.27)	(0.92)	(1.35)
γ_2	-7.412	0.334	0.732	0.779
	(-0.32)	(1.83)	(0.16)	(0.29)
χ^2	-	-	-	-
dof	0	0	0	0
p	-	-	-	-
GMM	$9.70e^{-14}$	$1.11e^{-15}$	$1.86e^{-14}$	$5.32e^{-17}$

Panel B: VIX-dependent non-monotonic discount factor

	1-month	2-month	3-month	6-month
γ_1	0.700	0.660	0.129	1.991
	(2.78)	(2.61)	(0.23)	(4.08)
γ_2	-47.470	-19.068	-2.143	-17.422
	(-1.26)	(-1.06)	(-0.30)	(-1.38)
γ_3	-8.153	-7.415	-9.720	-4.421
	(-2.21)	(-2.29)	(-0.43)	(-1.09)
χ^2	-	-	-	-
dof	0	0	0	0
p	-	-	-	-
GMM	$6.85e^{-13}$	$4.68e^{-14}$	$1.73e^{-14}$	$7.81e^{-14}$

Table 4 Summary Statistics for Realized Returns, Variances, Third Moments, Risk-Neutral Variances, and the NVIX

This table reports summary statistics for the data used in this study. Panel A reports means, standard deviations, and number of observations, N, for the realized S&P500 excess returns, variances, skewness, and the normalized VIX ($nvix$) across the four expirations: 1-, 2-, 3-, and 6-month. S&P500 excess returns between dates t and $t+T$, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, are calculated over the risk-free rate, which is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. Realized S&P500 variances, $v^2_{t,t+T}$, are calculated according to equation (24). Realized S&P500 third central moments, $s^3_{t,t+T}$, are calculated according to equation (25). Risk-neutral S&P500 variances, $var^{rnd}_t(R_{t,t+T})$, are calculated from the option-based risk-neutral density (equation (15)). $nvix_{t,t+T}$ is the normalized VIX, which is the VIX divided by its 1986-1995 average (to avoid look-ahead bias) and appropriately scaled for each expiration. Panel B reports correlations of realized excess returns, variances, third central moments, risk-neutral variance, and the $nvix$. Realized excess returns, variances, third central moments, risk-neutral variances, and the $nvix$ are for non-overlapping intervals. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Realized excess returns, variances, third central moments, and the $nvix$

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
mean	0.69%	1.27%	1.55%	3.79%
st. deviation	4.92%	7.05%	7.52%	10.86%
variance	0.24%	0.49%	0.56%	1.18%
skewness	-1.14	-1.61	-1.15	-0.81
3^{rd} central moment	-0.01%	-0.06%	-0.04%	-0.10%
$v^2_{t,t+T}$	1-month	2-month	3-month	6-month
mean	0.26%	0.53%	0.72%	1.53%
st. deviation	0.38%	0.72%	0.93%	1.54%
$s^3_{t,t+T}$	1-month	2-month	3-month	6-month
mean	-0.01%	-0.03%	-0.06%	-0.14%
st. deviation	0.12%	0.35%	0.25%	0.54%
$var^{rnd}_t(R_{t,t+T})$	1-month	2-month	3-month	3-month
mean	0.51%	1.04%	1.25%	2.67%
st. deviation	0.59%	1.12%	0.98%	1.52%
$nvix_{t,t+T}$	1-month	2-month	3-month	3-month
mean	1.42	1.42	1.00	1.27
st. deviation	1.79	1.98	1.05	0.94
N	323	147	83	52

Panel B: Correlations of realized excess returns, variances, third central moments, risk-neutral variances, and the nvix

	1-month				2-month				3-month				6-month			
	$R^e_{t,t+T}$	$v^2_{t,t+T}$	$s^3_{t,t+T}$	$var^{rnd}_t(R_{t,t+T})$	$R^e_{t,t+T}$	$v^2_{t,t+T}$	$s^3_{t,t+T}$	$var^{rnd}_t(R_{t,t+T})$	$R^e_{t,t+T}$	$v^2_{t,t+T}$	$s^3_{t,t+T}$	$var^{rnd}_t(R_{t,t+T})$	$R^e_{t,t+T}$	$v^2_{t,t+T}$	$s^3_{t,t+T}$	$var^{rnd}_t(R_{t,t+T})$
$v^2_{t,t+T}$	-0.19				-0.46				-0.53				-0.63			
$s^3_{t,t+T}$	0.57	-0.33			0.64	-0.60			0.63	-0.59			0.59	-0.68		
$var^{rnd}_t(R_{t,t+T})$	0.21	0.64	-0.11		0.17	0.41	0.11		0.03	0.61	-0.20		0.16	0.28	0.06	
$nvix_{t,t+T}$	0.17	0.67	-0.10	0.93	0.15	0.53	0.05	0.92	0.02	0.67	-0.13	0.90	0.18	0.38	0.07	0.84

Table 5 Risk-Return Regressions: Realized Returns and Variances

This table reports no-intercept regression results for equation (17). In Panel A, we regress realized S&P500 returns in excess of the risk-free rate, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, on realized S&P500 variances, $v^2_{t,t+T}$. Realized S&P500 variances are calculated according to equation (24). In Panel B, we regress, backward-looking fitted excess returns for S&P500, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - R^f_{t,t+T}$, on realized S&P500 variances. Backward-looking fitted S&P500 returns, $\widehat{R}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate. These regressions are reported in Table OA.1 of the Online Appendix. In Panel C, we regress, realized excess returns for S&P500 on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared. These regressions are reported Table OA.1 of the Online Appendix. In Panel D, we regress, backward-looking fitted excess returns for S&P500 on backward-looking fitted variances. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, and 4 lags for the 1-, 2-, and 3-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Realized excess returns on realized variances

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	-0.818	-2.128	-1.896	-1.012
	(-0.85)	(-2.03)	(-2.25)	(-1.00)
R^2	2.08%	15.66%	19.84%	16.09%
N	323	147	83	52

Panel B: Backward-looking fitted excess returns on realized variances

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.008	0.888	0.789	0.949
	(3.67)	(3.75)	(2.71)	(2.60)
R^2	-8.39%	-15.12%	-15.01%	-20.58%
N	321	146	82	50

Panel C: Realized excess returns on backward-looking fitted variances

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	2.616	1.817	1.152	2.583
	(1.89)	(1.48)	(0.77)	(2.22)
R^2	2.18%	-0.00%	-1.09%	3.71%
N	322	146	82	51

Panel D: Backward-looking fitted excess returns on backward-looking fitted variances

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.711	1.762	1.274	2.164
	(5.15)	(4.55)	(3.52)	(3.68)
R^2	0.41%	-3.42%	-11.22%	6.43%
N	321	146	82	50

Table 6 Summary Statistics for Forward-looking Expected Returns and Variances from Monotonic Discount Factors

This table reports summary statistics for the forward-looking expected returns and variances under the option-based physical measures. Panel A reports means, standard deviations, and number of observations, N, for the forward-looking S&P500 risk premia and variances across three expirations, 1-, 2-, 3-, and 6-month, based on the fixed-parameter monotonic stochastic discount factor (SDF) from equation (1). Panel B reports summary statistics for expected returns and variances under the physical measure based on the VIX-dependent parameter monotonic pricing kernel from equation (2). Forward-looking S&P500 risk premia between dates t and $t+T$, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, are over the risk-free rate, which is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R, t+T]$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, and variances, $var_t(R_{t,t+T})$, are calculated using the physical density, which is derived from the option-based risk-neutral density and the estimated discount factors (equation (15)). The estimation of the various pricing kernels is reported in Table 2. T is the number of days within each 1-, 2-, 3-, and 6-month interval ($T \approx 30, 60, 90, 180$). $\rho()$ is the correlation between forward-looking risk-premia, physical variances, and risk-neutral variances. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia and variances (fixed-parameter monotonic SDF)

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
mean	0.67%	1.28%	1.56%	3.80%
st. deviation	0.77%	1.36%	1.20%	2.11%
$var_t(R_{t,t+T})$	1-month	2-month	3-month	3-month
mean	0.42%	0.81%	0.96%	1.98%
st. deviation	0.44%	0.76%	0.69%	1.02%
$\rho\big(var_t(R_{t,t+T}), var^{rnd}_t(R_{t,t+T})\big)$	0.99	0.98	0.98	0.98
$\rho\big(\mathbb{E}_t[R^e_{t,t+T}], var_t(R_{t,t+T})\big)$	1-month	2-month	3-month	3-month
	0.99	0.98	0.98	0.98
N	323	147	83	52

Panel B: Forward-looking risk premia and variances (VIX-dependent monotonic SDF)

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
mean	0.69%	1.27%	1.55%	3.79%
st. deviation	2.61%	0.83%	0.64%	1.20%
$var_t(R_{t,t+T})$	1-month	2-month	3-month	3-month
mean	0.42%	0.82%	0.97%	1.99%
st. deviation	0.35%	0.87%	0.80%	1.19%
$\rho\big(var^{rnd}_t(R_{t,t+T}), var_t(R_{t,t+T})\big)$	0.96	0.93	0.99	0.99
$\rho\big(var_t(R_{t,t+T}), var^{rnd}_t(R_{t,t+T})\big)$	1-month	2-month	3-month	3-month
	0.85	0.96	0.89	0.88
N	323	147	83	52

Table 7 Risk-Return Regressions: Forward-looking Expected Returns and Variances from Monotonic Pricing Kernels with Fixed Parameters

This table reports no-intercept regression results for equation (17). In Panel A, we regress forward-looking risk premia for S&P500, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking S&P500 variances, $var_t(R_{t,t+T})$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, and variances are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the estimated pricing kernel of equation (1). The estimation of the pricing kernel is reported in Table 2, Panel A. In Panel B, we regress forward-looking risk premia for S&P500 on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared (Panel B, Table OA.1 of the Online Appendix). In Panel C, we regress forward-looking risk premia for S&P500 on realized S&P500 variances, $v^2_{t,t+T}$. Realized S&P500 variances are calculated according to equation (24). In Panel D, we regress fitted S&P500 returns in excess of the risk-free rate, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - Rr^f_{t,t+T}$, on forward-looking S&P500 variances. Backward-looking fitted S&P500 returns, $\widehat{R}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate (Panel A, Table OA.1 of the Online Appendix). In Panel E, we regress realized S&P500 returns in excess of the risk-free rate, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous, and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia on forward-looking variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	1.662	1.658	1.641	1.941
	(38.08)	(23.73)	(39.54)	(50.40)
R^2	98.46%	97.29%	96.68%	96.76%
N	323	147	83	52

Panel B: Forward-looking risk premia on backward-looking fitted variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	2.637	2.481	1.879	2.608
	(20.99)	(10.98)	(12.39)	(15.58)
R^2	92.37%	60.20%	76.17%	78.15%
N	322	146	82	51

Panel C: Forward-looking risk premia on realized variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.678	1.359	1.302	1.429
	(7.79)	(5.29)	(7.30)	(4.16)
R^2	38.50%	8.65%	21.52%	-48.52%
N	323	147	83	52

Panel D: Backward-looking fitted excess returns on forward-looking variances

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	1.095	1.062	1.299	1.560
	(5.67)	(5.27)	(5.45)	(3.74)
R^2	3.94%	-0.13%	0.43%	5.91%
N	321	146	82	50

Panel E: Realized excess returns on forward-looking variances

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	2.032	1.510	1.146	1.751
	(2.67)	(1.72)	(1.01)	(2.09)
R^2	4.49%	2.47%	-0.61%	0.86%
N	323	147	83	52

Table 8 Risk-Return Regressions: Forward-looking Expected Returns and Variances from Monotonic Pricing Kernels with VIX-dependent Parameters

This table reports no-intercept non-linear regression results for equation (19). In Panel A, we regress forward-looking risk premia for S&P500 $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking S&P500 variances, $var_t(R_{t,t+T})$, multiplied by $nvix_{t,t+T}$, which is the VIX divided by its 1986-1995 average (to avoid look-ahead bias) and appropriately scaled for each expiration. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, and variances are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the estimated pricing kernel of equation (2). The estimation for the pricing kernel is reported in Panel B of Table 2. In Panel B, we regress forward-looking risk premia for S&P500 on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$, multiplied by $nvix_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared (Panel B, Table OA.1 of the Online Appendix). In Panel C, we regress forward-looking risk premia for S&P500, on realized S&P500 variances, $v^2_{t,t+T}$ multiplied by $nvix$. Realized S&P500 variances are calculated according to equation (24). In Panel D, we regress fitted S&P500 returns in excess of the risk free rate, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - Rr^f_{t,t+T}$, on forward-looking S&P500 variances multiplied by $nvix_{t,t+T}$. Backward-looking fitted S&P500 returns, $\widehat{R}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate (Panel A, Table OA.1 of the Online Appendix). In Panel E, we regress realized S&P500 returns in excess of the risk-free rate, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances, multiplied by the $nvix_{t,t+T}$. In Panel F, we regress backward-looking fitted returns of backward-looking fitted variances, and in Panel G, we regress realized excess returns on realized variances. In both cases, the explanatory variables are multiplied by the $nvix$. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia on forward-looking variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	0.767	1.747	1.645	2.101
	(11.87)	(23.66)	(22.40)	(26.19)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	1.703	-0.468	-0.823	-0.799
	(18.96)	(-15.97)	(-25.65)	(-13.63)
R^2	98.49%	95.08%	89.30%	90.33%
N	323	147	83	52

Panel B: Forward-looking risk premia on backward-looking fitted variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	0.909	2.304	2.211	2.575
	(5.04)	(17.23)	(13.49)	(18.91)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	1.788	-0.172	-1.001	-0.351
	(13.69)	(-1.02)	(-8.07)	(-2.88)
R^2	98.20%	50.26%	48.21%	53.95%
N	322	146	82	51

Panel C: Forward-looking risk premia on realized variances

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	0.384	1.475	1.250	1.447
	(3.16)	(5.05)	(6.67)	(3.24)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	2.397	-0.469	-0.584	-0.573
	(12.20)	(-1.89)	(-3.89)	(-0.45)
R^2	89.14%	-72.46%	-202.19%	-439.70%
N	323	147	83	52

Panel D: Backward-looking fitted excess returns on forward-looking variances

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	2.069	1.947	1.697	1.776
	(3.65)	(3.85)	(4.77)	(1.54)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	-1.206	-1.346	-1.396	-0.562
	(-0.53)	(-0.65)	(-1.88)	(-0.29)
R^2	9.87%	5.76%	11.76%	5.89%
N	321	146	82	50

Panel E: Realized excess returns on forward-looking variances

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	0.998	1.251	1.596	1.574
	(1.15)	(1.03)	(1.74)	(1.76)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	0.929	0.144	-1.177	0.192
	(1.63)	(0.18)	(-1.32)	(0.30)
R^2	5.11%	2.63%	0.25%	1.02%
N	323	147	83	52

Panel F: Backward-looking fitted excess returns on backward-looking fitted variances

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	3.037	2.553	2.108	2.282
	(3.80)	(3.66)	(3.25)	(1.90)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	-1.012	-1.025	-1.560	-0.228
	(-1.25)	(-0.82)	(-1.76)	(-0.14)
R^2	3.10%	1.89%	-2.10%	6.64%
N	321	146	82	50

Panel G: Realized excess returns on realized variances

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	-3.986	-3.148	-2.586	-1.176
	(-2.51)	(-2.59)	(-2.59)	(-1.27)
$nvix^{\widetilde{\gamma}_3}$, $\widetilde{\gamma}_3$	-2.185	-0.745	-0.934	-0.585
	(-2.66)	(-1.14)	(-1.78)	(-1.51)
R^2	2.85%	5.07%	6.89%	-7.70%
N	323	147	83	52

Table 9 Summary Statistics for Forward-looking Expected Returns, Variances, and Third Central Moments from Non-monotonic Discount Factors

This table reports summary statistics for the forward-looking expected returns and variances under the option-based physical measures. Panel A reports means, standard deviations, and number of observations, N, for the forward-looking S&P500 risk premia, variances, and third central moments across three expirations, 1-, 2-, 3-, and 6-month, based on the fixed-parameter non-monotonic pricing kernel of equation (4). Panel B reports summary statistics for expected returns, variances, and third central moments under the physical measure based on the VIX-dependent non-monotonic pricing kernel from equation (6). Forward-looking S&P500 risk premia between dates t and $t+T$, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, are over the risk-free rate, which is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, variances, $var_t(R_{t,t+T})$, and third central moments, $m_{3,t}(R_{t,t+T}) = E_t[(R_{t,t+T} - E_t[R_{t,t+T}])^3]$, are calculated using the physical density from option prices (equation (15)). The estimation of the pricing kernels is reported in Table 3. T is the number of days within each 1-, 2-, 3-, and 6-month interval ($T \approx 30, 60, 90, 180$). Forward-looking risk premia and variances. $\rho()$ is the correlation between forward-looking risk-premia, physical variance, risk-neutral variances, and third central moments. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia, variances, and third central moments (fixed-parameter non-monotonic SDF)

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
mean	0.69%	1.27%	1.55%	3.79%
st. deviation	0.73%	1.05%	1.14%	1.96%
$var_t(R_{t,t+T})$	1-month	2-month	3-month	6-month
mean	0.34%	0.89%	1.05%	2.23%
st. deviation	0.27%	1.37%	0.83%	1.38%
$\rho\big(var_t(R_{t,t+T}), var^{rnd}_t(R_{t,t+T})\big)$	0.98	0.98	0.97	0.94
$m_{3,t}(R_{t,t+T})$	1-month	2-month	3-month	6-month
mean	-0.01%	-0.01%	-0.15%	-0.32%
st. deviation	0.01%	0.64%	0.18%	0.34%
$\rho\big(\mathbb{E}_t[R^e_{t,t+T}], var_t(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	0.92	0.60	0.99	0.98
$\rho\big(\mathbb{E}_t[R^e_{t,t+T}], m_{3,t}(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	-0.56	-0.28	-0.81	-0.82
$\rho\big(var_t(R_{t,t+T}), m_{3,t}(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	-0.44	-0.93	-0.85	-0.89
N	323	147	83	52

Panel B: Forward-looking risk premia, variances, and third central moments (VIX-dependent non-monotonic SDF)

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
mean	0.69%	1.27%	1.55%	3.79%
st. deviation	0.41%	0.71%	1.14%	1.11%
$var_t(R_{t,t+T})$	1-month	2-month	3-month	6-month
mean	0.36%	0.70%	0.87%	1.32%
st. deviation	0.62%	1.17%	1.10%	1.31%
$\rho\big(var_t(R_{t,t+T}), var_t^{rnd}(R_{t,t+T})\big)$	0.96	0.99	0.98	0.98
$m_{3,t}(R_{t,t+T})$	1-month	2-month	3-month	6-month
mean	-0.03%	-0.07%	-0.10%	-0.04%
st. deviation	0.10%	0.22%	0.22%	0.11%
$\rho\big(\mathbb{E}_t[R^e_{t,t+T}], var_t(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	-0.52	-0.55	-0.66	-0.38
$\rho\big(\mathbb{E}_t[R^e_{t,t+T}], m_{3,t}(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	0.39	0.38	0.52	0.31
$\rho\big(var_t(R_{t,t+T}), m_{3,t}(R_{t,t+T})\big)$	1-month	2-month	3-month	6-month
	-0.92	-0.96	-0.92	-0.92
N	323	147	83	52

Table 10 Risk-Return Regressions: Forward-looking Expected Returns, Variances, and Third Moments from Non-monotonic Pricing Kernels with Fixed Parameters

This table reports no-intercept regression results for equation (23). In Panel E, we regress forward-looking risk premia for S&P500, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking S&P500 , $var_t(R_{t,t+T})$, and third central moments, $m_{3,t}(R_{t,t+T})$. Forward-looking S&P500 expected returns, $\widetilde{\mathbb{E}}_t[R_{t,t+T}]$, variances, and third-moments are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the non-monotonic estimated pricing kernel of equation (4). The estimation of the pricing kernel is reported in Table 3, Panel A. In Panel B, we regress forward-looking risk premia for S&P500 on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$, and third moments, $\widehat{s}^3_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared. Backward-looking fitted variances are fitted values from regressing realized skewness on lag-realized skewness and the forward-looking third moment. These regressions are reported in Panels B and C, Table OA.1 of the Online Appendix. In Panel C, we regress forward-looking risk premia for S&P500 on realized S&P500 variances, $v^2_{t,t+T}$, and third central moment, $s^3_{t,t+T}$. Realized S&P500 variances are calculated according to equation (24). Realized S&P500 third central moments are calculated according to equation (25). In Panel D, we regress fitted S&P500 returns in excess of the risk-free rate, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances and third moments. Backward-looking fitted S&P500 returns, $\widehat{R}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate. These regressions are reported in Panel A, Table OA.1 of the Online Appendix. In Panel E, we regress realized S&P500 returns in excess of the risk-free rate, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances and third central moments. In Panel F, we regress backward-looking fitted returns of backward-looking fitted variances an third moments, and in Panel G, we regress realized excess returns on realized variances and third moments. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia on forward-looking variances and third moments

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	2.029	1.955	1.576	2.001
	(10.42)	(31.24)	(48.22)	(28.86)
$m_{3,t}(R_{t,t+T})$, $\widetilde{\gamma}_2$	-4.379	3.418	0.844	2.321
	(-1.18)	(24.65)	(4.48)	(5.29)
R^2	85.82%	92.95%	99.27%	97.59%
N	323	147	83	52

Panel B: Forward-looking risk premia on backward-looking fitted variances and third moments

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	2.295	2.400	1.158	3.029
	(32.75)	(24.71)	(5.61)	(16.67)
$\widehat{s}^3_{t,t+T}$, $\widetilde{\gamma}_2$	-4.934	-0.628	-9.580	5.727
	(-2.46)	(-1.68)	(-3.66)	(4.56)
R^2	89.79%	84.73%	86.07%	88.81%
N	322	146	82	51

Panel C: Forward-looking risk premia on realized variances and third moments

$\widetilde{\mathbb{E}}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.696	1.925	1.540	2.029
	(6.44)	(14.20)	(6.85)	(7.23)
$s^3_{t,t+T}$, $\widetilde{\gamma}_2$	0.871	2.381	1.925	3.535
	(1.76)	(4.35)	(1.59)	(1.95)
R^2	32.97%	53.91%	27.71%	-23.93%
N	323	147	83	52

Panel D: Backward-looking fitted excess returns on forward-looking variances and third moments

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	1.078	1.267	0.714	0.638
	(2.24)	(3.83)	(1.36)	(0.74)
$m_{3,t}(R_{t,t+T})$, $\widetilde{\gamma}_2$	-10.399	1.819	-2.663	-4.298
	(-0.93)	(2.66)	(-1.02)	(-1.10)
R^2	-0.36%	-2.08%	3.92%	14.76%
N	321	146	82	50

Panel E: Realized excess returns on forward-looking variances and third moments

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(r_{t,t+T})$, $\widetilde{\gamma}_1$	3.793	0.838	1.096	0.084
	(3.45)	(0.80)	(0.51)	(0.04)
$m_{3,t}(r_{t,t+T})$, $\widetilde{\gamma}_2$	30.458	-0.978	0.614	-8.767
	(1.70)	(-0.46)	(0.04)	(-0.88)
R^2	3.97%	4.78%	-0.68%	4.44%
N	323	147	83	52

Panel F: Backward-looking fitted excess returns on backward-looking variances and third moments

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.467	1.802	-0.735	2.654
	(4.27)	(4.20)	(-1.48)	(4.39)
$\widehat{s}^3_{t,t+T}$, $\widetilde{\gamma}_2$	-4.779	-2.350	-27.977	6.271
	(-1.61)	(-1.62)	(-3.85)	(1.25)
R^2	4.96%	-3.09%	10.79%	8.98%
N	321	146	82	50

Panel G: Realized excess returns on Realized variances and third moments

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	3.383	1.770	-0.556	4.268
	(3.63)	(1.34)	(-0.21)	(3.05)
$s^3_{t,t+T}$, $\widetilde{\gamma}_2$	15.010	2.834	-23.800	21.272
	(0.76)	(0.54)	(-0.84)	(1.59)
R^2	3.51%	0.35%	-0.14%	7.33%
N	323	147	83	52

Table 11 Risk-Return Regressions: Forward-looking Expected Returns, Variances, and Third Moments from Non-monotonic Pricing Kernels with VIX-dependent Parameters

This table reports no-intercept non-linear regression results for equation (22). In Panel A, we regress forward-looking risk premia for S&P500, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking S&P500 variances, $var_t(R_{t,t+T})$, and third central moments, $m_{3,t}(R_{t,t+T})$, both multiplied by $nvix_{t,t+T}$, which is the VIX divided by its 1986-1995 average (to avoid look-ahead bias) and appropriately scaled for each expiration. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, variances, and third moments are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the non-monotonic estimated pricing kernel of equation (6). The estimation of the pricing kernel is reported in Panel B of Table 3. In Panel B, we regress forward-looking risk premia for S&P500 on backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$, and third moments, $\widehat{s}^3_{t,t+T}$, both multiplied by $nvix_{t,t+T}$. Backward-looking fitted variances are fitted values from regressing realized variances on lag-realized variances and the VIX-squared. Backward-looking fitted variances are fitted values from regressing realized skewness on lag-realized skewness and the forward-looking third moment. These regressions are reported in Panels B and C, Table OA.1 of the Online Appendix. In Panel C, we regress forward-looking risk premia for S&P500 on realized S&P500 variances, $v^2_{t,t+T}$, and third central moments, $s^3_{t,t+T}$, where both dependent variables are multiplied by $nvix$. Realized S&P500 variances are calculated according to equation (24). Realized third central moments of S&P500 are calculated according to equation (25). In Panel D, we regress fitted S&P500 returns in excess of the risk free rate, $\widehat{R}^e_{t,t+T} = \widehat{R}_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances and third moments, both multiplied by $nvix_{t,t+T}$. Backward-looking fitted S&P500 returns, $\widehat{R}_{t,t+T}$, are fitted values from regressing realized returns on the price-dividend ratio, the dividend growth, and the risk-free rate. These regressions are reported in Panel A, Table OA.1 of the Online Appendix. In Panel E, we regress realized S&P500 returns in excess of the risk-free rate, $R^e_{t,t+T} = R_{t,t+T} - R^f_{t,t+T}$, on forward-looking S&P500 variances and third moments, both multiplied by the $nvix_{t,t+T}$. In Panel F, we regress backward-looking fitted excess returns on backward-looking fitted variances an third moments, and in Panel G, we regress realized excess returns on realized variances and third moments. In both cases, explanatory variables are multiplied by $nvix_{t,t+T}$. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia on forward-looking variances and third moments

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	0.323	0.399	0.049	3.903
	(3.72)	(1.96)	(1.32)	(13.92)
$m_{3,t}(R_{t,t+T})$, $\widetilde{\gamma}_2$	-182.010	-84.003	-7.535	-7.315
	(-14.65)	(-10.70)	(-5.49)	(-0.65)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-9.776	-8.573	-11.403	-3.462
	(-35.72)	(-14.39)	(-12.37)	(-14.88)
R^2	69.70%	70.26%	79.24%	29.67%
N	323	147	83	52

Panel B: Forward-looking risk premia on backward-looking fitted variances and third moments

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	3.385	1.419	0.274	3.365
	(15.28)	(5.58)	(1.70)	(23.13)
$\widehat{s}^3_{t,t+T}$, $\widetilde{\gamma}_2$	0.237	-8.727	-12.767	14.534
	(0.08)	(-5.27)	(-7.46)	(9.28)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-2.698	-1.625	-2.775	-1.242
	(-17.38)	(-7.57)	(-14.04)	(-22.81)
R^2	52.34%	41.30%	80.47%	42.17%
N	322	146	82	51

Panel C: Forward-looking risk premia on realized variances and third moments

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.922	1.806	1.094	2.292
	(7.36)	(7.60)	(5.72)	(7.30)
$s^3_{t,t+T}$, $\widetilde{\gamma}_2$	3.996	3.792	2.842	6.659
	(4.24)	(4.23)	(4.02)	(3.68)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-2.411	-2.030	-2.675	-1.695
	(-8.84)	(-6.84)	(-8.39)	(-7.31)
R^2	-117.62%	-109.38%	-44.17%	-336.15%
N	323	147	83	52

Panel D: Backward-looking fitted excess returns on forward-looking variances and third moments

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	0.578	1.641	0.531	2.246
	(0.20)	(3.66)	(0.76)	(0.56)
$m_{3,t}(R_{t,t+T})$, $\widetilde{\gamma}_2$	-140.87	-6.879	-10.907	-7.716
	(-1.19)	(-1.12)	(-5.35)	(-0.13)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-8.828	-1.789	-6.195	-0.985
	(-1.18)	(-0.77)	(-1.39)	(-0.18)
R^2	-3.91%	7.26%	-9.13%	-6.46%
N	321	146	82	50

Panel E: Realized excess returns on forward-looking variances and third moments

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$var_t(r_{t,t+T})$, $\widetilde{\gamma}_1$	1.478	0.602	0.495	0.687
	(0.88)	(0.25)	(0.32)	(0.40)
$m_{3,t}(r_{t,t+T})$, $\widetilde{\gamma}_2$	8.074	-11.278	-13.415	-105.30
	(0.95)	(-0.46)	(-1.18)	(-2.75)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	0.773	-0.646	-1.883	-1.873
	(0.81)	(-0.57)	(-2.01)	(-3.38)
R^2	5.63%	2.36%	1.12%	4.08%
N	323	147	83	52

Panel F: Backward-looking fitted excess returns on backward-looking fitted variances and third moments

$\widehat{R}^e_{t,t+T}$	1-month	2-month	3-month	6-month
$\widehat{v}^2_{t,t+T}$, $\widetilde{\gamma}_1$	3.090	2.554	-0.395	5.090
	(4.53)	(5.20)	(-0.433)	(5.02)
$\widehat{s}^3_{t,t+T}$, $\widetilde{\gamma}_2$	-10.697	3.624	-27.705	39.600
	(-2.79)	(1.13)	(-3.29)	(3.42)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-1.644	-1.014	-1.017	-1.110
	(1.49)	(-1.28)	(-1.83)	(-2.67)
R^2	12.39%	4.15%	19.52%	19.76%
N	321	146	82	50

Panel G: Realized excess returns on realized variances and third moments

$R^e_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t,t+T}$, $\widetilde{\gamma}_1$	1.078	0.752	0.994	0.171
	(1.06)	(0.92)	(1.25)	(3.96)
$s^3_{t,t+T}$, $\widetilde{\gamma}_2$	28.073	15.455	19.420	5.346
	(6.03)	(7.71)	(5.23)	(5.09)
$nvix^{\widetilde{\gamma}_3}_{t,t+T}$, $\widetilde{\gamma}_3$	-0.274	-0.336	0.356	5.528
	(-1.18)	(-1.06)	(0.88)	(12.18)
R^2	32.46%	37.82%	29.36%	48.74%
N	321	146	82	50

Table 12 Risk-Return Regressions: Standard Risk-Return Trade-off Regressions for the Non-Standard Discount Factors

This table reports linear regression results for equation (17), i.e., the standard risk-return trade-off, for the non-standard utility functions with fixed-parameter non-monotonicity (equation (4)) in Panel A, VIX-dependent monotonicity (equation (2)) in Panel B, and VIX-dependent non-monotonicity (equation (6)) in Panel C. In all Panels, we regress forward-looking risk premia for S&P500, $\mathbb{E}_t[R^e_{t,t+T}] = \mathbb{E}_t[R_{t,t+T}] - R^f_{t,t+T}$, on forward-looking S&P500 variances, $var_t(R_{t,t+T})$. Forward-looking S&P500 expected returns, $\mathbb{E}_t[R_{t,t+T}]$, and variances are calculated using the physical density from option prices (equation (15)), which is derived from the risk-neutral density and the three pricing kernels of equations (2), (4), and (6). Panel D reports results from no-intercept non-linear regressions of equation (19) with forward-looking risk premia and variances from the non-monotonic VIX-dependent pricking kernel of equation (6). The estimation of the non-standard pricing kernels is reported in Table 2, Panel B and Table 3. The risk-free rate is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. All variables are contemporaneous and all regressions are for non-overlapping intervals. t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2 lags for the 1-, 2-, 3-, and 6-month expirations, respectively. N is the number of observations. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Forward-looking risk premia and variances from fixed-parameter non-monotonic SDF

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	2.203	0.748	1.435	1.613
	(15.84)	(2.19)	(65.78)	(27.81)
R^2	84.74%	22.36%	98.79%	93.85%
N	323	147	83	52

Panel B: Forward-looking risk premia and variances from VIX-dependent monotonic SDF

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var(R_{t,t+T})$, $\widetilde{\gamma}_1$	3.551	1.214	1.233	1.635
	(3.36)	(13.68)	(7.33)	(11.20)
R^2	59.57%	83.40%	34.30%	24.15%
N	323	147	83	52

Panel C: Forward-looking risk premia and variances from VIX-dependent non-monotonic SDF

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	0.229	0.228	0.262	1.284
	(2.24)	(2.15)	(1.96)	(3.27)
R^2	-47.96%	-55.88%	-39.86%	-342.09%
N	323	147	83	52

Panel D: Forward-looking risk premia and variances from VIX-dependent non-monotonic SDF

$\mathbb{E}_t[R^e_{t,t+T}]$	1-month	2-month	3-month	6-month
$var_t(R_{t,t+T})$, $\widetilde{\gamma}_1$	4.033	3.047	0.697	4.095
	(10.02)	(8.04)	(5.07)	(11.15)
$nvix_{t,t+T}$, $\widetilde{\gamma}_3$	-5.595	-5.267	-7.272	-3.307
	(-19.08)	(-15.39)	(-15.49)	(-15.31)
R^2	35.03%	50.63%	67.50%	27.48%
N	323	147	83	52

The Risk-Return Trade-off Puzzle: Backward- versus Forward-Looking Expected Returns, Online Appendix

Online Appendix A Supplemental Figures

Figure OA.1 Risk-Neutral Density Functions

This figure shows the average risk-neutral density function across different expirations with and without tail adjustment. For these plots, we average the risk-neutral density functions (with and without tail adjustments) in each expiration across dates. The tail adjustment is done by appending a type-I (two-parameter) Pareto distribution to the tails of the risk-neutral density. The two parameters of the left-tail (right-tail) Pareto distribution are identified by matching the Pareto distribution to the empirical risk-neutral distributions at 2% (98%) and 5% (95%). Details on the derivation of the risk-neutral distribution can be found in Online Appendix C. The sample is from January 1996 to December 2022 for 1-month expiration options, May 1998 to November 2022 for 2-month expiration, January 2002 to October 2022 for 3-month expiration, and June 1996 to June 2022 for 6-month expiration options.

Panel A: One-month expiration

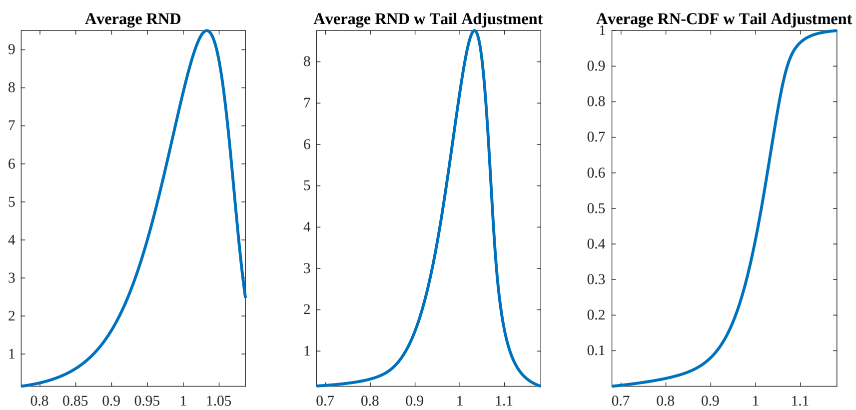

Panel B: Two-month expiration

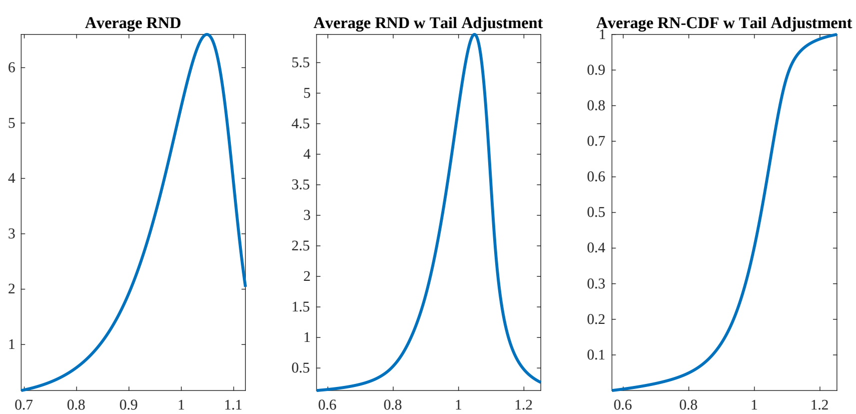

Panel C: Three-month expiration

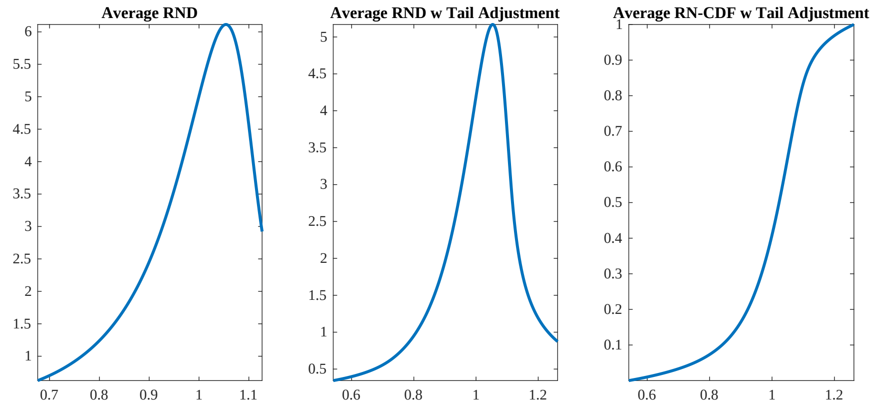

Panel D: Six-month expiration

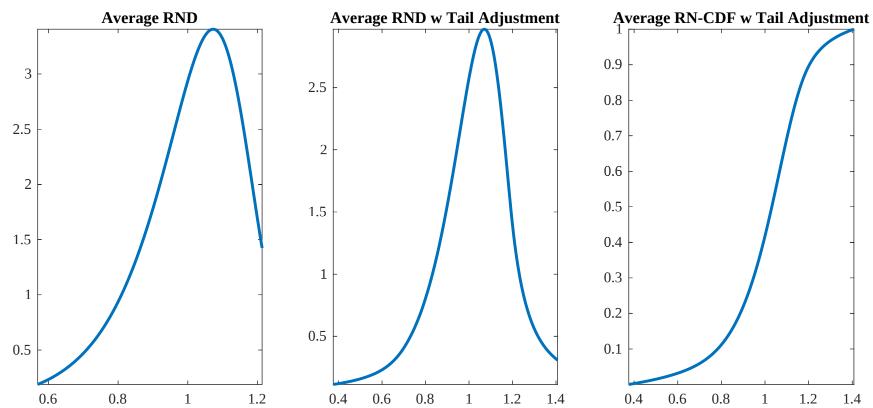

Online Appendix B Supplemental Tables

Table OA.1 Backward-looking Fitted Returns, Variances, and Third Moments

This table reports three sets of regression results for the backward-looking fitted returns, variances, and third central moments of the S&P500, which are used in the risk-return regressions. In Panel A, we derive the backward-looking fitted returns, $\widehat{R}_{t,t+T}$ by regressing realized returns, $R_{t,t+T}$, on the dividend-yield, $divp_t$, the lagged dividend growth, $\Delta div_{t-T,t}$, and the risk-free rate, $R^f_{t,t+T}$, which is set equal to the mean of the risk-neutral distribution, $R^f_{t,t+T} = E^{rnd}_t[R_{t,t+T}]$. In Panel B, we derive the backward-looking fitted variances, $\widehat{v}^2_{t,t+T}$, by regressing realized variances, $v^2_{t,t+T}$, on lag realized variances, $v^2_{t-T,t}$, and VIX-squared, $vix^2_{t,t+T}$. Realized variances for each expiration are derived according to equation (24). In Panel C, we derive the backward-looking fitted third central moment, $\widehat{s}^3_{t,t+T}$, by regressing realized third central moments, $s^3_{t,t+T}$, on lag realized third central moments, $s^3_{t-T,t}$, and the third central moment of the risk-neutral distribution $m^{rnd}_{3,t}(R_{t,t+T})$. Realized third central moments for each expiration are derived according to equation (25). The third central moment of the risk-neutral distribution is derived from equation (15). t-statistics in parentheses are corrected for heteroscedasticity and autocorrelation using Newey-West standard errors with 12, 6, 4, and 2, lags for the 1-, 2-, 3-, and 6-month expirations, respectively. All regressions are for non-overlapping intervals. Panel D reports means, variances, and number of observations, N, for fitted returns, variances, and third central moments. Realized stock market returns, dividends, and dividend-yield are from the CRSP Index files for the S&P500 index. The sample is from January 1996 to December 2022 for 1-month maturity, May 1998 to November 2022 for 2-month maturity, January 2002 to October 2022 for 3-month maturity, and June 1996 to June 2022 for 6-month maturity.

Panel A: Realized returns on dividend-yield, lagged dividend growth, and the risk-free rate

$R_{t,t+T} - 1$	1-month	2-month	3-month	6-month
$divp_t$	12.954	16.092	10.930	14.037
	(1.56)	(1.24)	(0.84)	(1.22)
$\Delta div_{t-T,t}$	-0.011	-0.018	0.029	-0.028
	(-1.15)	(-0.57)	(0.28)	(-0.25)
$R^f_{t,t+T} - 1$	-0.520	0.314	-0.719	-0.359
	(-0.66)	(0.35)	(-0.82)	(-0.21)
constant	-1.376%	-3.797%	-3.755%	-9.082%
	(-1.13)	(-0.96)	(-0.59)	(-0.847)
R^2	1.45%	2.25%	3.20%	8.59%
N	321	146	82	50

Panel B: Realized variances on lag realized variances and the VIX

$v^2_{t,t+T}$	1-month	2-month	3-month	6-month
$v^2_{t-T,t}$	0.209	0.462	-0.156	0.052
	(2.18)	(4.15)	(-0.68)	(0.27)
$vix^2_{t,t+T}$	0.391	0.083	0.729	0.300
	(3.64)	(0.81)	(3.04)	(1.39)
constant	0.053%	0.216%	0.082%	0.799%
	(3.16)	(4.75)	(1.10)	(3.28)
R^2	46.67%	31.44%	52.97%	12.79%
N	322	146	82	51

Panel C: Realized third central moments on lag realized third central moments and risk-neutral central moments

$s^3_{t,t+T}$	1-month	2-month	3-month	6-month
$s^3_{t-T,t}$	-0.288	-0.289	0.061	0.045
	(-2.38)	(-3.99)	(0.48)	(0.35)
$m^{rnd}_{3,t}(R_{t,t+T})$	0.445	0.033	0.164	-0.246
	(1.25)	(0.26)	(1.20)	(-0.83)
constant	0.014%	-0.040%	-0.023%	-0.248%
	(0.73)	(-1.64)	(-0.81)	(-1.40)
R^2	10.51%	7.70%	3.96%	1.40%
N	322	146	82	51

Panel D: Fitted returns, variances, and third central moments

$\widehat{R}_{t,t+T}$	1-month	2-month	3-month	6-month
mean	0.61%	1.24%	1.67%	3.78%
st. deviation	0.58%	1.05%	1.32%	3.08%
N	321	146	82	50

$\widehat{v}^2_{t,t+T}$	1-month	2-month	3-month	6-month
mean	0.27%	0.53%	0.73%	1.54%
st. deviation	0.26%	0.40%	0.68%	0.55%
N	322	146	82	51

$\widehat{s}^3_{t,t+T}$	1-month	2-month	3-month	6-month
mean	-0.01%	-0.03%	-0.06%	-0.14%
st. deviation	0.04%	0.09%	0.05%	0.06%
N	322	146	82	51

Online Appendix C   Derivation of the Risk-neutral Density

The derivation of the risk-neutral density (RND) follows the methodology in Figlewski (2010), Birru and Figlewski (2012), Linn et al. (2018), and Alexiou et al. (ming). The first step is to construct the implied volatility (IV) curve across strike prices. For strike prices (K) outside the $\pm 2\%$ range of the underlying spot price (S_t), we directly use IV's provided by OptionMetrics. For strike prices inside the $\pm 2\%$ range of the underlying spot price, we combine the IV's of OptionMetrics for puts (IV_p) and calls (IV_c) with the same strike price into a single point

$$IV\big(K \in (1 \pm 2\%)S_t\big) = \omega IV_p\big(K \in (1 \pm 2\%)S_t\big) + (1-\omega) IV_c\big(K \in (1 \pm 2\%)S_t\big),$$

where $\omega = (K_{max} - K)/(K_{max} - K_{min})$, and K_{max} and K_{min} are respectively the maximum and minimum strike prices in the $\pm 2\%$ moneyness range. As in Alexiou et al. (ming), this is done to avoid an artificial jump at the ATM region, which could arise from ATM puts potentially trading at higher IV relative to ATM calls.

Based on these IV points, we construct the IV curve by fitting a quintic spline with 1,000 moneyness nodes. Using the Black and Scholes (1973) formula, the IV curve is then converted into a curve of call option prices, $C_{t,t+T}(S_t R_{i,t}, S_t, \widetilde{r}_{f,t,t+T}, IV_{t,t+T}(K), div_{t,t+T})$, where $R_{i,t} = K_{i,t}/S_t$ is the moneyness (or gross return) for every strike price (node) i, $\widetilde{r}_{f,t,t+T}$ is the continuously-compounded risk-free rate (Federal Funds Rate), and $div_{t,t+T}$ is the continuously-compounded dividend yield from OptionMetrics. The risk-neutral density, $\widetilde{q}_{t,t+T}(SR) = d\widetilde{Q}_{t,t+T}(SR)/d(SR)$, is derived using the result in Breeden and Litzenberger (1978) where

$$\widetilde{q}_{t,t+T}(S_t R_{t,t+T}) = e^{T\ \widetilde{r}_{f,t,t+T}} \frac{\partial^2 C_{t,t+T}(S_t R_{t,t+T}, S_t, \widetilde{r}_{f,t,t+T}, IV_{t,t+T}(K), div_{t,+T})}{\partial (S_t R_{t,t+T})^2}.$$

The chain rule implies that the risk-neutral density for gross returns is

$$\widetilde{q}_{t,t+T}(R_{t,t+T}) = S_t e^{T\ \widetilde{r}_{f,t,t+T}} \frac{\partial^2 C_{t,t+T}(S_t R_{t,t+T}, S_t, \widetilde{r}_{f,t,t+T}, IV_{t,t+T}(K), div_{t,+T})}{\partial (S_t R_{t,t+T})^2}.$$

The second derivative above is calculated using a second-order centered difference approximation. Further, we can rescale the RND with the factor

$$\widehat{q}_{t,t+T}(R_{i,t}) = \frac{\widetilde{q}_{t,t+T}(R_{i,t})}{\sum_{we=1}^{1,000} \widetilde{q}_{t,t+T}(R_{i,t})(R_{i+1,t} - R_{i,t})}.$$

In this case, $\widehat{q}_{t,t+T}(R_{i,t})$ is a well-defined density function for gross returns since

$$\sum_{i=1}^{1,000} \widehat{q}_{t,t+T}(R_{i,t})(R_{i+1,t} - R_{i,t}) = \sum_{we=1}^{1,000} \frac{\widetilde{q}_{t,t+T}(R_{i,t})(R_{i+1,t} - R_{i,t})}{\sum_{we=1}^{1,000} \widetilde{q}_{t,t+T}(R_{i,t})(R_{i+1,t} - R_{i,t})} = 1.$$

The derived density is truncated at its tails because options for extreme values of the stock market index are dropped from the sample ($< \$3/8$) or the corresponding prices are zero. Hence, as in Figlewski (2010) and Linn et al. (2018), the final step in the derivation of the RND is to adjust its tails by appending heavy-tailed distributions. To this end, we assume that the left (l) and right (r) tails of the RND are given by two-parameter Pareto density functions, $f_l(R)$ and $f_r(R)$:

$$f_l(R) = \frac{\alpha_l (\lambda_l - R)^{-\alpha_l - 1}}{\lambda_l^{-\alpha_l}}, R \le \lambda_l, \quad f_r(R) = \frac{\alpha_r R^{-\alpha_r - 1}}{\lambda_r^{-\alpha_r}}, R \ge \lambda_r. \tag{OA.1}$$

The two parameters, α_l and λ_l (α_r and λ_r) are identified by solving a 2×2 system of non-linear equations where we set the left (right) Pareto density above equal to the values of the derived RND at the 2% and 5% (95% and 98%) percentiles. Finally, using the solutions for the parameters in the tail distributions, we extend the domain of moneyness by approximately 60% (30% in the left tail and 30% in the right tail) from 1,000 to 1,600 nodes and re-normalize the RND to obtain a well-defined density that integrates to one:

$$q_{t,t+T}(R_{i,t}) = \frac{\widehat{q}_{t,t+T}(R_{i,t})}{\sum_{i=1}^{1,400} \widehat{q}_{t,t+T}(R_{i,t})(R_{i+1,t} - R_{i,t})}.$$

Online Appendix D Proofs: Pricing Kernels and Log-normal Densities

This section derives the resulting distributions after combining linear or quadratic pricing kernels with a log-normal random variable. The log-normal probability density function is proportional to

$$\frac{1}{R} Exp\Big[\frac{(lnR - r^f)^2}{2\sigma^2}\Big].$$

The constants r^f and σ^2 are the location and scale parameters of the distribution. Based on the probability density function above for log-normal random variables, we derive the distributions corresponding to the linear and quadratic pricing kernels.

Linear Pricing Kernel and Log-normal Distribution

$$\frac{1}{R} Exp\Big[\gamma_1 lnR - \frac{(lnR - r^f)^2}{2\sigma^2}\Big] \propto \frac{1}{R} Exp\Big[-\frac{\big(lnR - (r^f + \gamma_1\sigma^2)\big)^2}{2\sigma^2}\Big].$$

Hence, $\mathbb{E}[lnR] = r^f + \gamma_1\sigma^2$ and $var\big(lnR\big) = \sigma^2$.

Quadratic Pricing Kernel and Log-normal Distribution

$$\frac{1}{R} Exp\Big[\gamma_1 lnR + \gamma_2 ln^2 R - \frac{(lnR - r^f)^2}{2\sigma^2}\Big] \propto \frac{1}{R} Exp\Big[-\frac{\big(lnR - \frac{r^f + \gamma_1\sigma^2}{1 - 2\gamma_2\sigma^2}\big)^2}{2\frac{\sigma^2}{1 - 2\gamma_2\sigma^2}}\Big].$$

Hence, $\mathbb{E}[lnR] = (r^f + \gamma_1\sigma^2)/(1 - 2\gamma_2\sigma^2)$ and $var\big(lnR\big) = \sigma^2/(1 - 2\gamma_2\sigma^2)$.