Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options*

Eben Lazarus

MIT Sloan

NOVEMBER 2022

Abstract

I present evidence from index options that the price of risk over the value of the S&P 500 increases as the investment horizon becomes shorter. I show first how these risk prices may be estimated from options data, by translating the risk-neutral probabilities implied by options prices into physical probabilities that must provide unbiased forecasts of the terminal outcome. The risk price can be interpreted as the marginal investor's effective risk aversion, and estimating this value over different option-expiration horizons for the S&P, I find that risk aversion is reliably higher for near-term outcomes than for longer-term outcomes. It is difficult to reconcile these findings with leading macro–finance models, and I discuss necessary conditions for any such rational model to produce such a pattern. Models with dynamically inconsistent risk preferences, however, are capable of straightforwardly producing the findings presented here, and I discuss possible specifications of such models and their applicability to related results from previous literature.

*Contact: elazarus@mit.edu. I am grateful to Ned Augenblick, Ravi Bansal, Michael Bauer (discussant), Jules van Binsbergen, John Campbell, Eduardo Dávila, Ian Dew-Becker, Emmanuel Farhi, Stefano Giglio, Niels Gormsen, David Laibson, Matthew Rabin, Martin Schmalz, Gill Segal, Jim Stock, Jessica Wachter, and participants at the University of Wisconsin Junior Finance Conference and the San Francisco Fed Conference on Advances in Financial Research for comments, guidance, and relevant discussions.

1. Introduction

How do people assess risky outcomes at different horizons? This question is central to understanding intertemporal choice in the face of risk, and it has accordingly received much attention in the recent finance literature. Several recent papers, beginning with van Binsbergen, Brandt, and Koijen (2012), have argued that the term structure of equity returns is downward-sloping, with claims to near-term dividends (short-term *dividend strips*) exhibiting larger risk premia than longer-horizon claims on average.[1] Such a finding may seem intuitive in light of the observed risk premium on value stocks, which tend to have shorter-duration cash flows than growth stocks (Campbell and Vuolteenaho, 2004; Lettau and Wachter, 2007). But a downward-sloping equity term structure in fact runs counter to the predictions of many leading equilibrium asset-pricing models, and a subset of the recent literature has challenged the empirical finding on the grounds of measurement-error and sample-selection issues.[2]

In this paper, I contribute to the evidence on risk pricing at different horizons by focusing specifically on digital (or binary) options over the market index value at short to medium horizons. I show how the market's effective risk aversion over the terminal index value may be estimated at varying horizons using these option prices, by translating the risk-neutral probabilities implied by options prices into physical probabilities that by definition provide unbiased forecasts of the terminal outcome. Then conducting such estimation using S&P 500 index options, I find evidence consistent with a downward-sloping term structure of risk prices, as the market's effective risk aversion is reliably higher for near-term outcomes than for longer-term outcomes. In particular, a statistic interpretable as relative risk aversion is estimated to be around 15 at a one-week horizon, but it decreases essentially monotonically to around 3 at a 12-week horizon.

The evidence I present extends previous findings along two dimensions. First, and most importantly, I show that the declining term structure of risk prices for binary options provides additional information on the source of the declining risk premium for dividend strips found in previous literature. The risk premium for a dividend strip at a given maturity depends on both risk preferences (roughly, the "price" of risk) and the data-generating process for consumption and dividends (roughly, the "quantity" of risk). For example, Hasler and Marfè (2016) show that a rare-disasters model extended to allow for recoveries following a disaster is capable of producing a downward-sloping term structure

[1] In addition to van Binsbergen, Brandt, and Koijen (2012), see, among others, Binsbergen, Hueskes, Koijen, and Vrugt (2013), Binsbergen and Koijen (2017), Gormsen (2021), and Weber (2018).

[2] See Boguth, Carlson, Fisher, and Simutin (2022), Bansal et al. (2021), and Song (2018).

of risk premia on dividend strips.[3] Meanwhile, by considering binary options over the index value as in this paper, my finding of a downward-sloping term structure of risk prices is more difficult to rationalize by appealing to features of the data-generating process alone. Intuitively, considering binary options allows me to fix the riskiness of outcomes across horizons on at least one dimension (in my case, the percent deviation in the index across the two possible index-value outcomes).

Second, the use of index options data leads me to consider risk pricing at the short end of the term structure. Previous literature has tended to focus on either medium- to long-maturity pricing (see Footnote 1) or very long-term pricing, as in the case of Giglio, Maggiori, and Stroebel (2015).[4] In addition to providing new evidence for the shorter end of the term structure, these short-term options have the further advantage that they allow for risk-price estimation using the returns on buy-and-hold claims. This is in contrast to much of the literature examining longer maturities, where holding-period returns are used given the short time span of available observations. As discussed by Boguth, Carlson, Fisher, and Simutin (2022), this leads to possibly biased inference in the presence of measurement error, which is mitigated by using buy-and-hold returns, as done here. Further, I can account directly for the possibility of measurement error by instrumenting my main estimation equation with lagged risk-neutral probabilities, which I show does not affect the estimated results.[5]

Summarizing my estimation procedure in more detail (but without the full formal apparatus built up in Section 2), the key steps are as follows:

1. Options allow for bets over the future asset price, and thus the prices of these bets can be transformed into a probability distribution over the price at expiration using standard techniques.

2. This probability distribution (referred to as the *risk-neutral distribution*) can be transformed into a set of *conditional* probabilities over binary outcomes — in particular, the probability that the index return over a fixed horizon T will be A conditional on it being either A or B — as in Augenblick and Lazarus (2022).

3. These conditional risk-neutral probabilities (for now, π_t^* at date t) are in general

[3]The intuition can be expressed with a simple example following Binsbergen et al. (2013). Consider a disaster-and-recovery process such that if a disaster strikes consumption and dividends in period $t+1$, then those values are expected to fully recover in $t+2$. As of time t, the one-period dividend strip is fully exposed to the $t+1$ disaster risk, whereas the two-period strip is not, and it accordingly commands a lower premium.

[4]One exception is Dew-Becker, Giglio, Le, and Rodriguez (2017), who document a declining term structure of variance-risk prices over the first few monthly maturity horizons.

[5]See Boguth, Carlson, Fisher, and Simutin (2011) for an argument in favor of similar (though differently specified) instrumentation in a related context.

distorted relative to the true *physical* probabilities (π_t) over the binary outcome given the presence of risk aversion. In particular, it can be shown that there is a one-to-one relationship between π_t^* and π_t that depends only on the marginal investor's effective risk aversion over the binary (A vs. B) outcome.

4. The value π_t is unobserved, but it must be an unbiased forecast of the terminal outcome by definition. I can thus use the terminal outcomes themselves to estimate the degree of risk aversion embedded in π_t^* such that the implied π_t value has zero average forecast error for that terminal outcome (at all possible values of π_t^*).

5. Varying $T - t$ (or fixing $t = 0$ and varying T) allows for such estimation at varying horizons, holding fixed the binary return outcome.

The unbiased-forecast condition in step 4 can be estimated straightforwardly using the generalized method of moments.

This estimation procedure is similar in spirit to that of Hansen and Jagannathan (1991), who show how risk premia may be related to the variance of the SDF process. I am essentially using the estimated option risk premia at different horizons to obtain an estimate of the slope of the SDF across return states. Unlike in their setting, I obtain a point estimate rather than a bound, which is an advantage of this option-pricing setting since options allow me to condition on the terminal value itself. Bliss and Panigirtzoglou (2004) conduct similar estimation in an option-pricing setting closer to mine, focusing on average risk-aversion estimates and using stronger parametric testing assumptions, and see also Aït-Sahalia and Lo (1998) and Rosenberg and Engle (2002) for estimation using yet-stronger assumptions on the underlying data-generating processes.

The proposed estimation method also resembles what some previous literature — see Lichtenstein, Fischhoff, and Phillips (1977) for an early reference — has referred to as "calibration," where individual forecast rationality is tested (in cases where probabilistic forecasts are directly observable) by testing whether, for example, a given event happens 30 percent of the time on average when a given forecaster gives that outcome a 30 percent ex-ante probability of occurring. As Augenblick and Rabin (2021) note, this estimation is extremely inefficient given that it occurs pointwise across the entire distribution of ex-ante probability forecasts. In my case, I effectively integrate across the entire distribution of ex-ante forecasts to obtain a single relevant moment condition at each horizon, which again yields an estimate of the SDF slope across return states at different horizons (rather than a test of rationality per se).

The question then becomes how to interpret the finding that risk aversion increases as the investment horizon becomes shorter. I first derive a necessary condition under which such a finding would arise naturally in a fully rational framework: it must be the case that risk aversion over the terminal return outcome decreases with marginal utility. Since most standard models feature contemporaneous increases in marginal utility and risk aversion during bad times, it is difficult (though not impossible) to reconcile my empirical findings with such models. By contrast, models with dynamically inconsistent risk preferences are capable of straightforwardly producing the findings I present. I argue that such models may be interpreted as reduced-form versions of models in which loss-averse agents narrowly frame the outcomes of individual gambles (see Barberis and Huang, 2008; Rabin and Weizsäcker, 2009), and that my empirical evidence points in favor of narrower framing at shorter horizons.

The remainder of the paper is organized as follows. Section 2 introduces the theoretical framework and derives moment conditions for estimating risk prices over index returns at varying horizons. Section 3 then describes the data and presents my main empirical results, while Section 4 discusses their interpretation in the context of various theoretical frameworks. Section 5 concludes. The Appendix contains additional technical material.

2. Framework for Estimation

I first lay out the theoretical framework used to guide the estimation procedure. The setting, presented in Section 2.1, is a slightly simplified version of the framework presented in Augenblick and Lazarus (2022); I relegate additional technical detail to Appendix A.1. Section 2.2 then discusses the estimation procedure.

2.1. Theoretical Setting

Consider a discrete-time economy with time $t \in \{0, 1, 2, \ldots\}$. Denote the ex-dividend value of the market index by V_t. I will be concerned with the realization of uncertainty over the value V_T for some option expiration date T (or, more generally, some set of option expiration dates $\{T\}$). Denote the set of possible terminal index values (or some subset thereof) by $\mathcal{V}_T \equiv \{v_1, v_2, \ldots, v_J\}$,[6] ordered such that $v_1 < v_2 < \ldots < v_J$, and consider two arbitrary adjacent members of this set, v_j, v_{j+1}. Denoting the *physical* or *objective* probability measure by $\mathbb{P}$,[7] denote the time-t probability of the terminal index value being equal to v_j,

[6]I treat the set of possible index values as discrete to avoid additional technicalities and notational complication, but the analysis could be extended without loss to accommodate continuous state spaces.

[7]More formally, I assume a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t \in \mathbb{N}}$.

conditional on being either v_j or v_{j+1}, by

$$\pi_{t,j} \equiv \mathbb{P}_t(V_T = v_j \mid V_T \in \{v_j, v_{j+1}\}). \tag{1}$$

Under the absence of arbitrage, there exists a strictly positive *stochastic discount factor (SDF)* process $\{M_t\}$ such that the time-t price of a claim to an arbitrary state-contingent payoff X_T is given by $\mathbb{E}_t[(M_T/M_t)X_T]$, where $\mathbb{E}$ is the expectation under $\mathbb{P}$. This implies the existence of a *risk-neutral measure* $\mathbb{P}^*$ such that the time-t price of the same payoff X_T can equivalently be written as $\mathbb{E}_t^*[X_T]/R_{t,T}^f$, where $\mathbb{E}^*$ is the expectation under $\mathbb{P}^*$ and $R_{t,T}^f$ is the $(T-t)$-period gross risk-free rate at date t. Define the risk-neutral analogue to the conditional probability in (1) as

$$\pi_{t,j}^* \equiv \mathbb{P}_t^*(V_T = v_j \mid V_T \in \{v_j, v_{j+1}\}). \tag{2}$$

This risk-neutral probability can be measured from the set of option prices on date t expiring on date T with different strikes K, using standard results as first presented by Breeden and Litzenberger (1978).[8] It can be seen that

$$\pi_{t,j}^* = \frac{\mathbb{E}_t[M_T/M_t \mid V_T = v_j]}{\mathbb{E}_t[M_T/M_t \mid V_T \in \{v_j, v_{j+1}\}]} \, \pi_{t,j}. \tag{3}$$

An odds-ratio transformation of this equation yields

$$\frac{\pi_{t,j}^*}{1-\pi_{t,j}^*} = \phi_{t,T,j} \, \frac{\pi_{t,j}}{1-\pi_{t,j}}, \tag{4}$$

$$\text{where} \quad \phi_{t,T,j} \equiv \frac{\mathbb{E}_t[M_T/M_t \mid V_T = v_j]}{\mathbb{E}_t[M_T/M_t \mid V_T = v_{j+1}]}.$$

The values $\{\phi_{t,T,j}\}$ will be the objects of interest in the empirical exploration below. Intuitively, $\phi_{t,T,j}$ represents the price of risk over the bad (low-index-value) state relative to the good state, as encoded in the slope of the SDF across the two states. In the case in which a representative agent faces the consumption process $\{C_t\}$ and has time-separable consumption utility and rational expectations, this value becomes $\phi_{t,T,j} = \mathbb{E}_t[U'(C_T) \mid V_T = v_j]/\mathbb{E}_t[U'(C_T) \mid V_T = v_{j+1}]$. With the additional restriction that the representative agent in fact has (indirect) utility over time-T wealth, with wealth equal to the market index value, then Augenblick and Lazarus (2022, Proposition 6) show that relative risk aversion

[8] For concreteness, a European call option on the index with strike price K (where $K \in \mathcal{K} \subseteq \mathbb{R}_+$) has payoff $X_{T,K} = \max\{V_T - K, 0\}$. As above, see Appendix A.1 for further details.

$\gamma_{t,T,j} \equiv -v_j U''(v_j)/U'(v_j)$ is given to a first order around v_j by

$$\gamma_{t,T,j} = \frac{\phi_{t,T,j} - 1}{(v_{j+1} - v_j)/v_j}. \tag{5}$$

Relative risk aversion is proportional in this case to $\phi_{t,T,j} - 1$, as is intuitive given that this gives the percent decrease in marginal utility obtained by moving from the bad index outcome to the good outcome. To obtain relative risk aversion, this change in marginal utility must be normalized by the "consumption" increase in moving from the bad state to the good state, as in the denominator of (5).

I will in particular be interested in how the price of risk $\phi_{t,T,j}$ changes on average with the horizon $T-t$. Augenblick and Lazarus (2022) make the assumption, referred to there as *conditional transition independence*, that $\phi_{t,T,j}$ is constant over t for fixed j and T. I do not make such an assumption, and in fact one interpretation of the results below is that they provide direct empirical tests of that assumption. I do, however, make two simplifying assumptions as follows.

ASSUMPTION 1 (***Scale Independence***)**.** For arbitrary pairs of index values (v_j, v_{j+1}) and (v_k, v_{k+1}) for terminal date T, if $v_{j+1}/v_j = v_k/v_{k+1}$, then $\phi_{t,T,j} = \phi_{t,T,k}$. ‖

Given (5), this assumption can be interpreted as a constant-relative-risk-aversion assumption.[9] I will in fact assume that the set of possible terminal values $\mathcal{V}_T$ is equally (proportionally) spaced so that $v_{j+1}/v_j = v_k/v_{k+1}$ for all j, k, as can be enforced by construction in the empirical implementation below. The value $\phi_{t,T,j}$ can accordingly be written as $\phi_{t,T}$. The following assumption then allows for additional simplification.

ASSUMPTION 2 (***Horizon Dependence***)**.** The value $\phi_{t,T}$ depends only on the horizon $T-t$ for all dates and terminal dates, and accordingly write this vaue as ϕ_{T-t}. ‖

These two assumptions are made largely for the purposes of notational simplification and so that I can pool estimates across state pairs and expiration dates below. (I could, for example, instead simply define $\phi_{T-t} \equiv \mathbb{E}[\phi_{t,T,j}]$, where the average is taken over all dates and state pairs, and make appropriate stationarity assumptions so that the GMM procedure below provides a meaningful estimate of such an average.)

[9] The approximate constancy of the scale of the equity risk premium in the U.S. (see, for example, Caballero, Farhi, and Gourinchas, 2017, and Martin, 2017) can be taken as indirect evidence in favor of such scale independence, and see Campbell (2018) for further discussion.

2.2. Estimation of Horizon-Dependent Risk Pricing

2.2.1. Moment Condition

As discussed after equation (2), option prices allow for essentially direct observation of risk-neutral probabilities (up to issues of measurement error, to be discussed below). But physical probabilities are unobservable, yielding a continuum of possible solutions to equation (4), the mapping between physical and risk-neutral probabilities. A rewriting of that equation, however, makes clear how ϕ_{T-t} may nonetheless be estimated consistently in the data. First, rearrange that equation (applying Assumptions 1–2) as

$$\pi_{t,j} = \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})}. \tag{6}$$

Since $\pi_{t,j} = \mathbb{E}_t[\, \mathbb{1}\{V_T = v_j\} \mid V_T \in \{v_j, v_{j+1}\}]$ by definition, we have

$$\mathbb{E}_t\left[\mathbb{1}\{V_T = v_j\} - \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})} \;\middle|\; V_T \in \{v_j, v_{j+1}\} \right] = 0. \tag{7}$$

Note that the random variable $\mathbb{1}\{V_T = v_j\}$ is observable as of date T, as it simply indexes whether the terminal index value is equal to v_j. Thus every value in (7) is in principle observable aside from π_{T-t}, so applying the law of iterated expectations to this equation yields a nonlinear moment condition for ϕ_{T-t} that can be estimated using the generalized method of moments (GMM).

Economically, what this moment condition entails is estimation of the price-of-risk parameter needed to reconcile the ex-ante market forecast of the terminal outcome (as in $\pi^*_{t,j}$) with the average outcomes themselves. One can see from (6) that in general, given the ordering $v_j < v_{j+1}$ so that ϕ_{T-t} is likely greater than 1 in the presence of risk aversion,[10] it is the case that $\pi^*_{t,j} > \pi_{t,j}$; given my labeling, $\pi^*_{t,j}$ is the risk-neutral probability for the "bad" state, which in general will be higher than the true physical probability of that state occurring given the insurance value embedded in a contract that pays off in a bad state of the world. That insurance value is indexed exactly by the value ϕ_{T-t}, and the moment condition implied by (7) simply uses the insight that one can infer that insurance value by setting ϕ_{T-t} such that the ex-post forecast errors between the implied $\pi_{t,j}$ and

[10] The "risk-aversion puzzle" documented by Jackwerth (2000) possibly confounds this general economic intuition, though see Chabi-Yo, Garcia, and Renault (2008) and Linn, Shive, and Shumway (2018) for evidence that this finding is not robust to proper conditioning on other date-t variables. Following these latter papers, it will be the case in my reported results that I do not observe any such risk-aversion puzzle, as my point estimates all indicate that $\phi_{T-t} > 1$.

the observed $\mathbb{1}\{V_T = v_j\}$ must be mean-zero. Since these forecast errors are orthogonal to date-t information by definition, there are no endogeneity-related concerns.[11] The data then provides quasi-experimental variation in the horizon $T-t$, as I can observe repeated iterations of (7) for different forecast horizons with no ex-ante distinction between the data-generating processes on these different dates, allowing for estimation of ϕ_{T-t} across different horizons.

2.2.2. Measurement Error and an Implementable Orthogonality Condition

One possible concern with such estimation is the likelihood of price measurement error affecting the measured risk-neutral probabilities in (7) given, for example, market microstructure noise. Unlike in the case of Augenblick and Lazarus (2022), the GMM framework used here allows for me to account directly for this noise without needing to estimate its magnitude separately. First, assume that the observed conditional risk-neutral belief $\widehat{\pi}^*_{t,j}$ is measured with additive error with respect to the true value $\pi^*_{t,j}$ used in (7):

$$\widehat{\pi}^*_{t,j} = \pi^*_{t,j} + \epsilon_{t,j}, \tag{8}$$

where $\mathbb{E}[\epsilon_{t+k,j}\, \pi^*_{t+k',j} \mid V_T \in \{v_j, v_{j+1}\}] = 0$ for all k, k', and $\epsilon_{t,j}$ follows an MA(q) for some value q. (This is a slight relaxation of the assumptions used for the noise process in Augenblick and Lazarus, 2022, where it is effectively assumed that $\epsilon_{t,j}$ follows an MA(0).) It is shown in Appendix A.2 that the observed analogue of the second term in (7) is given by

$$\frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1-\widehat{\pi}^*_{t,j})} = \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1-\pi^*_{t,j})} + \epsilon_{t,j} + \underbrace{\mathcal{O}\Big(\big(\epsilon_{t,j} + (\phi_{T-t}-1)\big)^2\Big)}_{\text{higher-order terms}} \tag{9}$$

as $\epsilon_{t,j} \to 0$ and $\phi_{T-t} \to 1$,[12] where the latter limit $\phi_{T-t} = 1$ corresponds to the case of risk-neutrality as seen in (4).

Thus equation (7) can be re-expressed up to higher-order terms as

$$\mathbb{E}_t\left[\mathbb{1}\{V_T = v_j\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1-\widehat{\pi}^*_{t,j})} \;\middle|\; V_T \in \{v_j, v_{j+1}\}\right] = -\epsilon_{t,j}. \tag{10}$$

[11] If, however, the market systematically mis-forecasts future outcomes in one particular direction, then this will affect the estimated ϕ_{T-t} values, as these departures from rationality are embedded in the SDF sequence by construction. I discuss this possibility after presenting the empirical results below.

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

The risk-neutral probabilities used on the left side of this equation are now the observable values (inclusive of noise, unlike the ideal values used in (4)). Since $\epsilon_{t,j}$ is assumed to follow an MA(q), I can then form a set of unconditional moments by instrumenting using lagged values of $\widehat{\pi}^*_{t,j}$, for any lags greater than q.

That is, defining the N-dimensional instrument vector

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

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

$$\mathbb{E}\left[\left(\mathbb{1}\{V_T = v_j\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} \right) Z_{t,j} \,\middle|\, V_T \in \{v_j, v_{j+1}\} \right] = 0,$$

or, using the definition of the conditional expectation,

$$\mathbb{E}\left[\left(\mathbb{1}\{V_T = v_j\} - \frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} \mathbb{1}\Big\{ V_T \in \{v_j, v_{j+1}\} \Big\} \right) Z_{t,j} \right] = 0. \tag{11}$$

This unconditional moment restriction is now amenable to empirical estimation. Note from (10) that the instrument $Z_t = 1$ would in fact yield unbiased estimates of the parameter ϕ_{T-t}. But it is advantageous to use lagged-value instruments given both (a) the efficiency gains from doing so (Hayashi and Sims, 1983; Hansen, 1985), and (b) the fact that they allow for overidentification tests for the joint hypothesis that (7) and (8) are correctly specified.

The moment condition (11) can then be estimated over many expiration dates T, horizons $T - t$, and state pairs j. In particular, to account explicitly for the latter, denote a date-T-observable M-dimensional data vector by $X_{t,T}$, and define the function $h\colon \mathbb{R}^M \times$

$\mathbb{R} \to \mathbb{R}^{(J-1)\cdot N}$ as

$$h(X_{t,T}, \phi_{T-t}) = \begin{pmatrix} \left(\mathbb{1}\{V_T = v_1\} - \frac{\widehat{\pi}^*_{t,1}}{\widehat{\pi}^*_{t,1}+\phi_{T-t}(1-\widehat{\pi}^*_{t,1})}\,\mathbb{1}\left\{V_T \in \{v_1, v_2\}\right\}\right)\overrightarrow{1} \\ \vdots \\ \left(\mathbb{1}\{V_T = v_{J-1}\} - \frac{\widehat{\pi}^*_{t,J-1}}{\widehat{\pi}^*_{t,J-1}+\phi_{T-t}(1-\widehat{\pi}^*_{t,J-1})}\,\mathbb{1}\left\{V_T \in \{v_{J-1}, v_J\}\right\}\right)\overrightarrow{1} \end{pmatrix},$$

where $\overrightarrow{1}$ is an N-vector of ones. Define the full instrument vector $Z_t = (Z'_{t,1}, \cdots, Z'_{t,J-1})'$. Then the moment condition from which I can estimate ϕ_{T-t} for a given horizon $T-t$ is

$$\mathbb{E}\left[\, h(X_{t,T}, \phi_{T-t})'\, Z_t \,\right] = 0. \tag{12}$$

The expectation is taken over all pairs $t = \tau_1, T = \tau_2$ such that $\tau_2 - \tau_1 = \kappa$, in order to identify ϕ_κ. One can then stack the moment condition in (12) for values of $\kappa = 1, 2, \ldots$, to obtain horizon-dependent risk-price estimates, as I do in the estimation below.

3. Empirical Estimation and Main Results

3.1. Data Description

As in Augenblick and Lazarus (2022), I use S&P 500 index options data from the Option-Metrics database, which lists end-of-day bid and ask prices for European options on the index value over the sample January 1996–December 2018. This yields data for 5,537 trading dates and 991 expiration dates. I drop any options with bid prices of zero (or less than zero), with Black–Scholes implied volatility of greater than 100 percent, or with greater than 12 weeks to maturity (given the relative lack of observations and statistical power beyond this maturity), and calculate each option's end-of-day price as the midpoint between its bid and ask prices.

For each observed expiration date T and associated initial option trading date 0, I define the relevant (sub)set of possible terminal index values for the remainder of the empirical analysis as

$$\mathcal{V}_T = (V_0 R^f_{0,T}) \exp\Big(\big\{[-0.10, -0.08), [-0.08, -0.06), \ldots, [0.06, 0.08), [0.08, 0.10)\big\}\Big). \tag{13}$$

In words, state v_1 is said to be realized when the gross index-price appreciation, in excess of the risk-free rate $R^f_{0,T}$, is between $\exp(-0.1)$ and $\exp(-0.08)$, or equivalently when the log excess return is between -10% and -8%, and analogously for $v_2, \ldots, v_{10}$. Note that the states are equally spaced, as required by Assumption 1. Further, I exclude all terminal states more than 10% out of the money (where moneyness is relative to a zero excess return) in either direction, in order to avoid excessive measurement error in the tails of the distribution, but this does not require any assumption that the full set of possible terminal states is itself finite.[13]

I again follow the procedure in Augenblick and Lazarus (2022), due originally to Malz (2014) and building from the results of Breeden and Litzenberger (1978) discussed above after equation (2), to obtain observed risk-neutral probabilities $\widehat{\pi}^*_{t,j}$ (where the terminal date T is suppressed for simplicity) from the relevant option-price cross-sections; see Appendix A.3 for detail. Note again that these risk-neutral probabilities are conditional on state j or $j+1$ being realized.[14] I can also observe the realization of $\mathbb{1}\{V_T = v_j\}$ for all pairs T, j directly from the S&P 500 index price data for days on which the option settles at the end of the trading day, and I manually collect the settlement values for A.M.-settled options for this calculation from the Chicago Board Options Exchange website.

I exclude any T, j pairs for which $V_T \notin \{v_j, v_{j+1}\}$, since their contribution to the sample version of the moment condition in (11) is identically zero. This leaves 549 observations (tuples (t, T, j)) at the one-day horizon, which declines monotonically to 222 observations at the 60-day horizon (equivalently, the 12-week horizon), which motivates my focus on 1- to 12-week horizons as above.

3.2. Estimation and Results

I conduct estimation using GMM for sample counterparts of the moment condition (12). I make one further simplification relative to Assumptions 1–2 in this estimation: while I use daily data in constructing my sample moments,[15] I restrict ϕ_{T-t} to be fixed by weeks to expiration. Thus, for $T-t$ in days, I set $\phi_1 = \phi_2 = \ldots = \phi_5$, and so on. In reporting results below, I in fact refer to ϕ_1 as the one-week-horizon estimated value, and so on through ϕ_{12} for 12 weeks.

I apply this restriction for two main reasons. First, it greatly reduces the computational

[13] Such an assumption is made, for example, by Ross (2015); see Borovička, Hansen, and Scheinkman (2016) for a critical discussion.

[14] I also keep only conditional risk-neutral beliefs $\widehat{\pi}^*_{t,j}$ for which the non-conditional terminal-state beliefs satisfy $\widehat{\pi}^*_t(V_T = v_j) + \widehat{\pi}^*_t(V_T = v_j) \geqslant 5\%$, in order to reduce measurement error.

[15] That is, I have 60 moment conditions of the form (12), one for each horizon $T-t$ in days.

burden in estimation to decrease the number of estimated parameters by a factor of five (especially with respect to the bootstrap procedure used for inference), without sacrificing the essential economic insights of the estimation. Second, it allows me to obtain overidentifying restrictions even in the case where I use just one instrument (one lagged observed risk-neutral probability) for each moment equation, as is the case in my baseline estimation below.[16]

In my baseline estimation, I use the five-day-lagged observed risk-neutral probability $\widehat{\pi}^*_{t-5,j}$ as an instrument in the moment equation for $\widehat{\pi}^*_{t,j}$; following the discussion in Section 2.2.2, this is equivalent to assuming an MA(4) measurement-noise process and setting $\overline{q} = q + 1 = 5$, and I can directly test this assumption by examining the J-statistic arising from GMM estimation. I have experimented as well with a wide range of different lagged values as instruments (as well as the case in which no instrument is used); in all these cases, the estimates exhibit essentially identical patterns to those shown in the baseline case in this section, with risk prices declining significantly by horizon, and those results are available upon request. Details of my estimation procedure, as well as my method of inference for the purpose of constructing confidence intervals, can be found in Appendix A.4.

Figure 1 shows the main estimation results for ϕ_κ by week, along with pointwise 95% confidence intervals. I show the raw values $\widehat{\phi}_\kappa$, though the "price of risk" should in fact be thought of as $\widehat{\phi}_\kappa - 1$, given that $\phi_\kappa = 1$ corresponds to the case of risk neutrality and rational expectations, as can be seen in (4). This case is shown with a dotted line in the figure.

We can see immediately a clear downward-sloping pattern of risk-price estimates as the horizon increases, at least until about the six-week point, beyond which the values are insignificantly different from 1. To give a sense of the economic magnitudes implied by these estimates, note from equation (5) that we can interpret $(\phi_\kappa - 1) \times 50$ as the coefficient of relative risk aversion for an agent with utility over the index level itself, where the multiple 50 arises because I am using two-percentage-point bins as in (13) so that $(v_{j+1} - v_j)/v_j = 0.02$. This yields point estimates for relative risk aversion of 14.7 at the one-week horizon (95 percent confidence interval $[10.4, 18.9]$), 9.5 at the two-week horizon (CI $[6.4, 12.6]$), down to 3.4 at the 12-week horizon (CI $[0.1, 6.6]$). The J-statistic resulting from this estimation has a p-value of 0.30, indicating little evidence against the joint hypothesis that (7) and (8) are correctly specified, with the noise process in (8)

[16]In addition, as shown by Plagborg-Møller (2016, Chapter 3), given that we have *a priori* reasons to believe that the prices of risk are smooth across horizons, there may be mean-squared-error benefits to imposing this smoothness, as I do here in a particularly simple way by pooling estimates across days by week to expiration.

Figure 1: Estimates of Risk Prices by Horizon

Estimation by Two-Step GMM with Five-Day-Lag Instrument

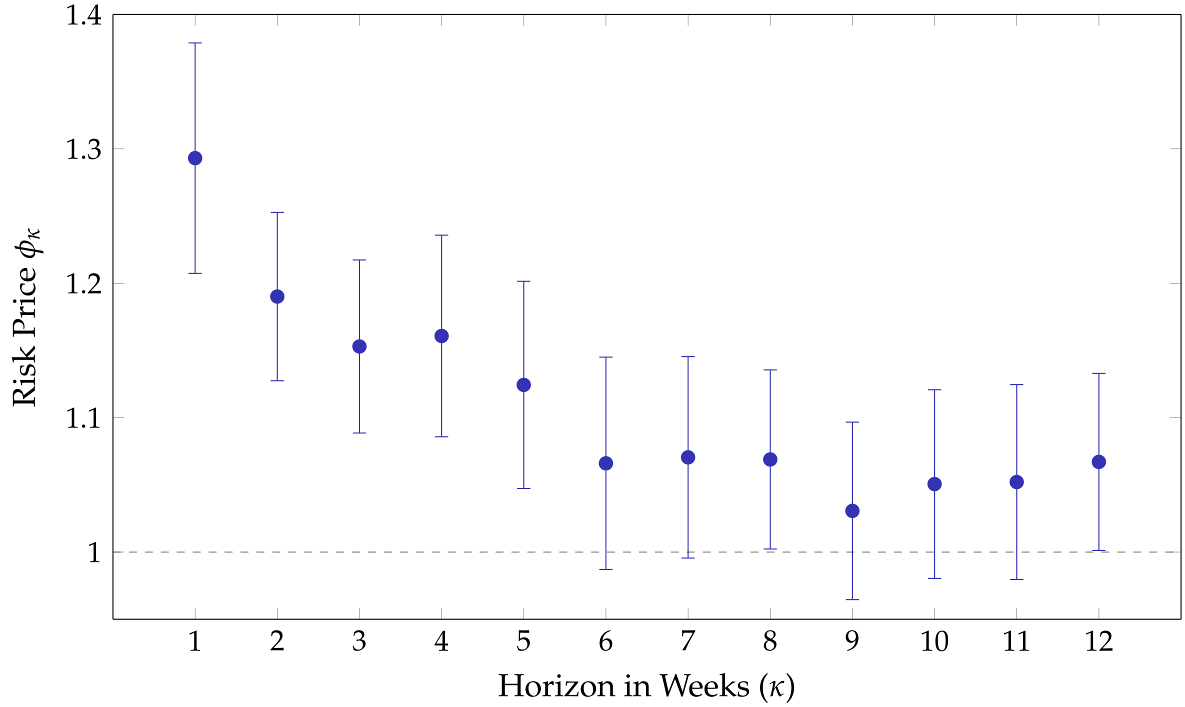

Notes: Point estimates are constructed using two-step GMM, using the five-day-lagged observation as an instrument, on the sample counterparts of the moment conditions in equation (12) in order to minimize forecast error. The price of risk parameter is constrained to be equal for all days within a given weekly horizon to expiration. Error bars show 95% confidence intervals, constructed using procedure in Appendix A.4. See that appendix for further details.

following an MA(q), $0 \leqslant q \leqslant 4$, as assumed in my estimation.[17]

In order to more formally assess whether the downward slope by horizon in Figure 1 is in fact a statistically robust finding across horizons, I estimate the following regression:

$$\widehat{\phi}_\kappa = \alpha + \beta\kappa + \varepsilon_\kappa. \tag{14}$$

That is, I run a regression of the estimated risk prices on a constant and a "horizon trend" κ, testing whether the trend β is significantly different than zero. For inference I use the block bootstrap discussed in Appendix A.4: I re-estimate ϕ_κ on 500 redrawn bootstrap samples,

[17] Further, none of the J-statistics across the robustness checks I have conducted (available upon request) reject that joint hypothesis at any conventional significance level. For simplicity, these p-values are constructed using asymptotic χ^2 critical values as originally developed by Hansen (1982) and applied in Hansen and Singleton (1982), and which may be problematic in time-series contexts; see, e.g., Hall and Horowitz (1996), Sun and Kim (2012), Lazarus, Lewis, and Stock (2021). As documented in those papers, however, *over*rejection tends to be the issue when using χ^2 critical values, so the fact that I am not rejecting the null suggests that this concern is not binding in the current setting.

rerun the regression (14) within each of these samples, and then calculate the distribution of the statistic $\widehat{\beta}^* - \widehat{\beta}$, where $\widehat{\beta}^*$ is the bootstrap estimate for β and $\widehat{\beta}$ is the estimate in the original sample. Denoting by $q^*(\cdot)$ the quantile function of the bootstrap distribution of $\widehat{\beta}^* - \widehat{\beta}$, I then calculate the 95% confidence interval as $[\widehat{\beta} - q^*(0.975), \widehat{\beta} - q^*(0.025)]$.

Conducting the above procedure, I obtain

$$\widehat{\beta} = -0.018,$$
$$95 \text{ percent CI } [-0.041, -0.007].$$

That is, the risk prices are estimated to decrease by roughly 0.02 by week to expiration (or, in terms of relative risk aversion, roughly 1 per week), and this is estimated as significantly different from zero in a two-sided 95 percent test. I thus conclude that risk pricing is horizon-dependent, with greater prices of risk at short horizons, and the remainder of the paper discusses how to interpret this finding.

4. Interpretation of Empirical Results

4.1. Rationalizing the Data in a Standard Framework

I begin by asking what features a standard, rational-expectations asset-pricing framework would require in order to generate the finding documented in Section 3. For this purpose, it is useful to consider a simple example. Assume a two-period horizon, $T = 2$, and two possible terminal index values V_2, denoted L, H, where $L < H$, with equal ex-ante probabilities. The terminal index values are not perfect proxies for the representative agent's marginal utility and the SDF. In particular, assume that there are two possible SDF realizations in each state, denoted as follows:

$$M_2 = \begin{Bmatrix} \mathfrak{a}_L & \text{with date-0 probability } 0.5 \\ \mathfrak{b}_L & \text{with date-0 probability } 0.5 \end{Bmatrix} \quad \text{if } V_2 = L,$$

$$M_2 = \begin{Bmatrix} \mathfrak{a}_H & \text{with date-0 probability } 0.5 \\ \mathfrak{b}_H & \text{with date-0 probability } 0.5 \end{Bmatrix} \quad \text{if } V_2 = H.$$

Normalize $M_0 = M_1 = 1$; this normalization is without loss of generality for determining conditional risk-neutral probabilities, since these depend only on $\mathbb{E}_t[M_2 \mid V_2 = L]/\mathbb{E}_t[M_2 \mid V_2 = H]$, as can be seen in (4).

The information and probability structure is illustrated graphically in Figure 2. I assume

Figure 2: Resolution of Uncertainty in Example

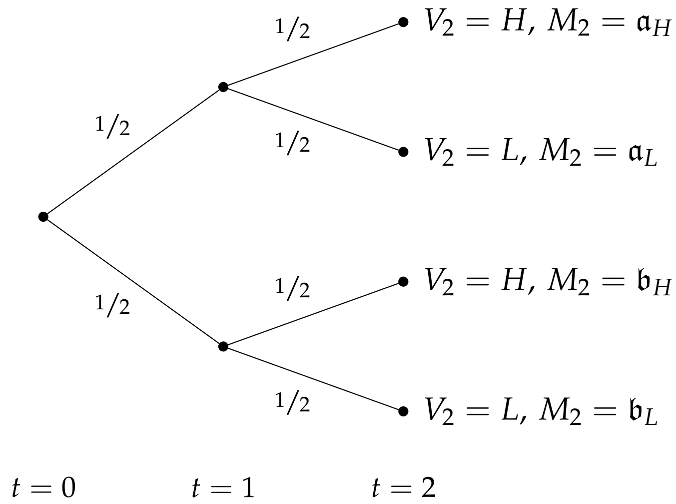

that as of date 1, there is no information revealed about whether the terminal index value will be L or H, so that those probabilities stay at 0.5,[18] but it is revealed what the SDF realization will be conditional on each state being realized: the representative agent learns either that $\mathfrak{a}_j$ will be realized if $V_2 = j$ is realized for $j = L, H$ (i.e., $\mathfrak{a}_L$ in state L or $\mathfrak{a}_H$ in H), or that $\mathfrak{b}_j$ will be realized if j is realized.

Since state L is the bad wealth state, set $\mathfrak{a}_L > \mathfrak{a}_H$, $\mathfrak{b}_L > \mathfrak{b}_H$, and we can label $\mathfrak{b}_j$ as the bad marginal-utility state in either case (corresponding to, e.g., high stochastic volatility, low long-run growth, lower surplus consumption in a model with habit formation), so that $\mathfrak{b}_j > \mathfrak{a}_j$ for $j = L, H$.

I can now ask under what conditions it would be the case that $\phi_0 < \mathbb{E}_0[\phi_1]$, where, as originally introduced in Section 2.1,

$$\phi_t \equiv \frac{\mathbb{E}_t[M_2 \mid V_T = L]}{\mathbb{E}_t[M_2 \mid V_T = H]}.$$

As can be seen from Figure 2, the condition $\phi_0 < \mathbb{E}_0[\phi_1]$ can be stated as

$$\phi_0 = \frac{\mathfrak{a}_L + \mathfrak{b}_L}{\mathfrak{a}_H + \mathfrak{b}_H} < \frac{1}{2}\frac{\mathfrak{a}_L}{\mathfrak{a}_H} + \frac{1}{2}\frac{\mathfrak{b}_L}{\mathfrak{b}_H} = \mathbb{E}_0[\phi_1],$$

[18] This is an appropriate assumption given the finding in the previous section that risk aversion is higher on average for each fixed value $\pi_{t,j}$ across the entire set of possible probabilities, though see the discussion below in Footnote 19.

which, under the normalization $\mathfrak{b}_j > \mathfrak{a}_j$ for $j = L, H$, yields, after simplification,

$$\frac{\mathfrak{a}_L}{\mathfrak{a}_H} > \frac{\mathfrak{b}_L}{\mathfrak{b}_H}.$$

That is, denoting ϕ_1 at the upper node for $t = 1$ by $\phi_{\mathfrak{a}}$, and similarly for the lower node by $\phi_{\mathfrak{b}}$, it must be the case that $\phi_{\mathfrak{a}} > \phi_{\mathfrak{b}}$.

Economically, what this requires is that risk aversion over the terminal index value be *higher* when the agent receives information that times are *good* in the sense that the part of marginal utility unrelated to the index return is expected to be low. The long-horizon gamble on the good-state outcome must be a good hedge (relative to the bad-state outcome) against bad intermediate marginal-utility news in order to generate a negative risk premium for the holding-period return on this gamble. (Note that this must be the case given that we observe increasing risk premia for such a gamble, held to maturity, in the data as the horizon becomes shorter.) That is, when an agent receives bad news about marginal utility, it must be the case that the relative price of the good-state gamble increases, which occurs when ϕ_t decreases.[19] Preliminary exploration indicates that this intuition can be shown to hold in more general cases, and this will be an interesting topic of future work on this subject.

The above condition is in general not met in leading representative-agent asset-pricing models. As an example, Figure 3 plots average risk prices by monthly (not weekly) horizon, as defined in the previous section, in the simulated long-run risks model of Bansal and Yaron (2004). Using code from Pohl, Schmedders, and Wilms (2018), I solve the model numerically using the calibration of Bansal and Yaron (2004, Case II) with stochastic volatility.[20] I then calculate average risk prices by months to expiration over 2,000,000 years of simulated monthly data. The risk prices are increasing very slightly by horizon, though not enough to be visible given the scale of the y-axis (set to be equivalent to the scale of Figure 1 for comparison).

Intuitively, when times are bad in the model, in the sense that marginal utility is high — i.e., when either stochastic volatility is high or expected long-run consumption growth is low — risk aversion over the terminal index value increases very slightly, violating the requirement derived above for the declining term structure of risk prices. See Gormsen (2021) for further discussion, as the requirement he derives to rationalize the cyclical

[19]This can also occur when the bad-state probability π_t decreases relative to π_{t-1}, which was assumed away in this example, but again this requires bad news about marginal-utility growth to be concurrent with good news about the return state.

[20]The solution uses projection methods and approximates expectations via quadrature in order to capture higher-order effects. See Pohl, Schmedders, and Wilms (2018) for further details.

Figure 3: Risk Prices by Horizon in the Long-Run Risks Model

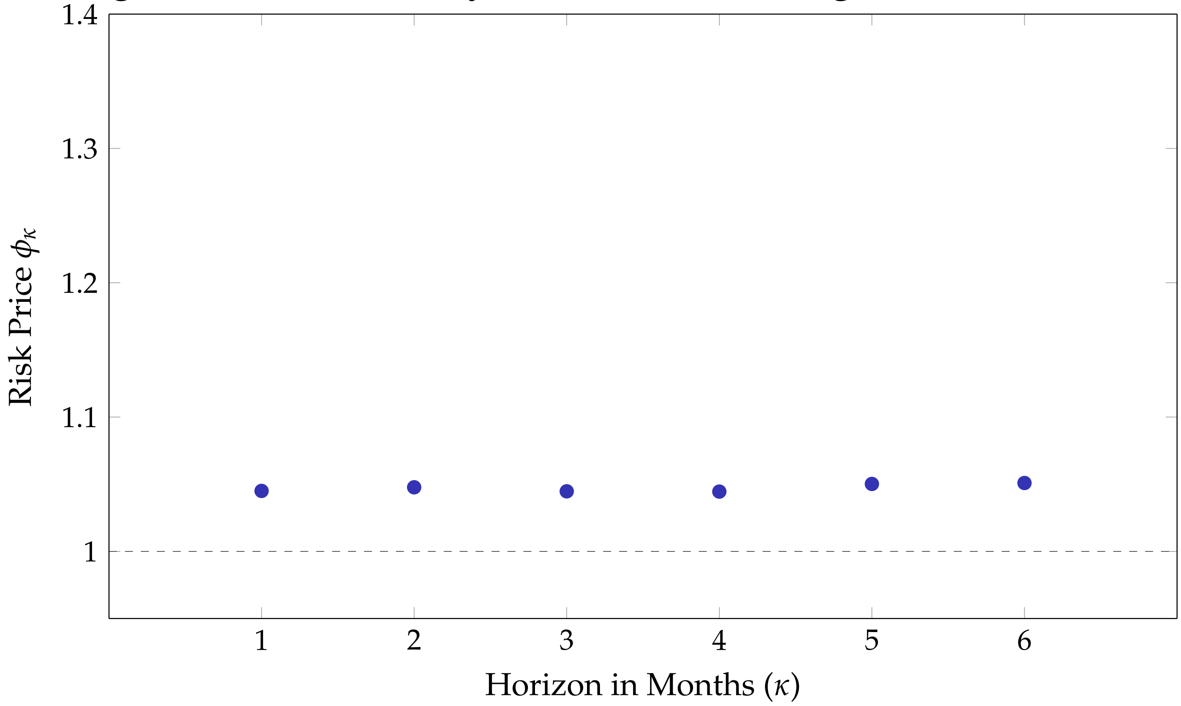

Notes: Risk prices are calculated as averages over 2,000,000 years of simulated monthly data, following the formula in equation (4). The model and calibration are as given in Bansal and Yaron (2004), where I use their Case II calibration. I solve the model numerically using the projection method of Pohl, Schmedders, and Wilms (2018).

variation in the equity term structure is quite similar to the requirement derived here.

4.2. Dynamically Inconsistent Risk Preferences

Departing from the standard representative-agent frameworks above, I can now ask what set of alternative assumptions could generate the patterns observed in the data. While there are likely to be many such frameworks, perhaps the simplest way of explaining the declining term structure of risk prices would be to take the declining relative-risk-aversion estimates at face value and assume that agents have different risk preferences over outcomes at different horizons. This is in fact exactly the tack taken by Eisenbach and Schmalz (2016) and Andries, Eisenbach, and Schmalz (2019), who motivate their approach by appealing both to experimental evidence and the previous asset-market evidence on downward-sloping risk premia.[21]

[21] Eisenbach and Schmalz (2016) include a review of experimental evidence of preference reversals as the horizon to uncertainty resolution decreases, as individuals seem to become more risk-averse or anxious about a lottery (or, in other settings, they get stage-fright on the day of a performance and regret having

I briefly present a version of the model considered by Andries, Eisenbach, and Schmalz (2019), who generalize Epstein–Zin (1989) preferences to include horizon-dependent risk aversion. Utility V_t is given by

$$V_t = \left[(1-\delta) C_t^{1-\frac{1}{\psi}} + \delta \mathbb{E}_t \left[\widetilde{V}_{t+1}^{1-\gamma_1} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma_1}} \right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{15}$$

where continuation utility $\widetilde{V}_{t+1}$ follows the recursion

$$\widetilde{V}_{t+1} = \left[(1-\delta) C_{t+1}^{1-\frac{1}{\psi}} + \delta \mathbb{E}_t \left[\widetilde{V}_{t+2}^{1-\gamma_2} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma_2}} \right]^{\frac{1}{1-\frac{1}{\psi}}}. \tag{16}$$

The case $\gamma_1 = \gamma_2$ is the usual Epstein–Zin (1989) case, with no dynamic inconsistency. When $\gamma_1 > \gamma_2$, however, risk aversion over near-term outcomes is greater than over distant-horizon outcomes. (This can be generalized to incorporate many different values over different horizons.) Andries, Eisenbach, and Schmalz (2019) show that regardless of the level of sophistication of the marginal (or representative) investor with respect to her dynamically inconsistent risk preferences, this leads to a declining term structure of equity risk premia.

The above specification is semi-reduced-form in the sense that it simply takes as given that risk preferences differ by horizon. But one way of rationalizing this framework in the context of pre-existing work on non-standard risk preferences may be to tie it to the literature on narrow framing and the equity premium begun by Benartzi and Thaler (1995). They propose that equity premia are higher than justified solely by the exposure of equity to consumption risk, because people frame lotteries narrowly (so they experience gains and losses with respect to equity returns themselves) and are loss-averse. See also Barberis and Huang (2008) for a more recent survey, as well as Rabin and Weizsäcker (2009) for a decision-making formalization of the earlier evidence and discussion of Tversky and Kahneman (1981).

One possible downside of the narrow-framing approach is that there are no clear guidelines as to what choice problems are narrowly framed.[22] For example, do individuals narrowly frame every individual-stock-level investment decision, over all horizons? The empirical results in the previous section suggest that this may not be the case: one inter-

volunteered to perform). See also Loewenstein (1996) for an earlier review across multiple domains.

[22]This relates to the literature discussing "free parameters" in behavioral models; see, for example, Wachter (2002) for a discussion.

pretation of that evidence is that near-term, salient outcomes are narrowly framed, which in combination with loss aversion causes effective short-horizon risk aversion to increase, at least for the overall equity index.[23] This empirical approach can accordingly be thought of as a disciplining mechanism for the specification of narrow framing, and perhaps leads toward models of dynamically inconsistent risk preferences as discussed above.

4.3. Preferences over the Timing of Resolution of Uncertainty

The above evidence that near-term outcomes command higher risk premia may seem to point in favor of a preference for late resolution of uncertainty, following the definition of Kreps and Porteus (1978). This is not quite the case; the near-term outcomes are both revealed and paid in the near term, whereas a test of preferences over the timing of the resolution of uncertainty would require a comparison of outcomes paying out at the same horizon, but with the payout value revealed early in some cases.[24]

There are nonetheless possible tests that do speak more directly to this preference. I have not yet implemented these tests in the data, so I relegate the details to Appendix A.5, but the intuition can be summarized briefly here. One can construct dynamic strategies that generate early-resolution lotteries with late payoffs, simply by reinvesting the proceeds of an early-resolving (and early-paying) option in a risk-free security that then pays off at the desired (late) horizon. If the risk-free rate is uncorrelated with the index (and therefore with the payoff of the early-resolving option), then this strategy is effectively as risky as a strategy without the risk-free reinvestment, allowing the results above to speak to the preference over the resolution of uncertainty. But in the case that the risk-free rate is correlated with the payoff of the early-resolving option — for purposes of intuition, assume the correlation is positive — then this increases the riskiness of the early-resolution/late-payoff strategy relative to the late-resolution/late-payoff strategy, if one maintains the same bins (which index the scale of the relative lottery payoffs) across option horizons as in (13). It is thus as of yet unclear whether the options data suggests a preference for the timing of uncertainty resolution in either direction.

[23] Note, however, that this is captured only at a high level and in reduced form in the specification of Andries, Eisenbach, and Schmalz (2019) presented above, given that their preferences are not themselves narrowly framed, and they are not loss-averse as they are continuously differentiable everywhere.

[24] Formally, Andries, Eisenbach, and Schmalz (2019) show that risk preferences can be dynamically inconsistent in the manner above, with $\gamma_1 > \gamma_2$ in (15)-(16), in such a way as to nonetheless not yield any clear prediction on the preference over the timing of resolution of uncertainty, as shown in their equation (5). Nonetheless, their Corollary 1 shows that this horizon-dependent risk aversion unambiguously lowers the timing premium relative to the benchmark case in which $\gamma_1 = \gamma_2$, even though the sign of the timing premium is ambiguous.

5. Conclusion

This paper presents evidence in favor of a declining term structure of risk prices with respect to gambles over small changes in the market index value over short to medium horizons; equivalently, it appears as if the market is more risk-averse with respect to short-horizon uncertainty over the index value than longer-horizon uncertainty. While I have discussed some classes of interpretations of the data, arguing here in favor of models with dynamically inconsistent risk preferences, further work remains to be done with respect to other classes of interpretations. It remains to be seen whether, for example, certain heterogeneous-agent models may be capable of rationalizing these findings.

The findings here may speak as well to the interpretation of the findings of Augenblick and Lazarus (2022), who find evidence against the rational-expectations assumption in the data when considering the volatility of the risk-neutral probability processes used here. The current paper has said little about the rationality of forecasts: I use the definitional unbiasedness property of physical probabilities to construct risk-price estimates, and those risk-price estimates can in theory incorporate both the effects of risk aversion and any average forecast errors for the marginal investor. But the fact that the risk-price estimates imply quite reasonable risk-aversion values (even at short horizons) seems to indicate that such forecasts are closed to unbiased on average,[25] though this of course does not preclude the excess volatility in conditional forecasts found by Augenblick and Lazarus (2022). Additional work remains to be done in understanding the two sets of results in a unified framework.

[25]Further, these reasonable risk-aversion estimates stand in contrast to the equity premium puzzle observed when considering the equity-index value itself, as documented by Mehra and Prescott (1985) and Hansen and Jagannathan (1991). But as discussed by Aït-Sahalia and Lo (2000) and Bliss and Panigirtzoglou (2004), risk-aversion estimates obtained from option-price-based forecasts are in fact often much more reasonable than the values obtained in the equity-premium-puzzle literature, so my results add further evidence in favor of this pattern, particularly in the middle of the index-return distribution.

Appendix: Additional Technical Material

A.1. Theoretical Framework: Technical Details

This appendix section presents the technical detail underlying the framework introduced in Section 2.1, largely following the setup in Augenblick and Lazarus (2022), which includes additional discussion.

I consider a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathbb{N}}$. A realization of the elementary state is denoted by $\omega \in \Omega$. I will be concerned with the ex-dividend value of the market index, $V_t \colon \Omega \to \mathbb{R}_+$, on some option expiration date T (or set of dates $\{T\}$); the subscript t will refer generally to $\mathcal{F}_t$-adapted processes. A European call option on the market index with strike price K has payoff $X_{T,K} = \max\{V_T - K, 0\}$, and denote its time-$t$ price as $q_{t,K}$. Assume without loss of generality that these option prices are observable for some set of strike prices $\mathcal{K} \subseteq \mathbb{R}_+$ beginning at date 0.

These option prices will be of interest for inferring a distribution over the change in value of the market index from 0 to T, or equivalently, fixing the first trading date 0 and $\mathcal{F}_0$, the value of the market index as of T. For notation, say that *index state* $v_j \in \overline{\mathcal{V}}_T \subset \mathbb{R}_+$ is realized for the market index as of date T if $V_T = v_j$, and I will consider an ordered subset $\mathcal{V}_T \subseteq \overline{\mathcal{V}}_T$, where $\mathcal{V}_T \equiv \{v_1, v_2, \ldots, v_J\}$, and $v_1 < v_2 < \ldots < v_J$. The measure $\mathbb{P} \colon \mathcal{F} \to [0,1]$ governs the *objective* or *physical* probabilities of these index states. The time-t objective probability that the index realizes state v_j at date T is

$$\mathbb{P}_t(V_T = v_j) = \sum_{\omega\,:\,V_T(\omega)=v_j} \mathbb{P}_t(\omega), \tag{A.1}$$

where $\mathbb{P}_t(\cdot) \equiv \mathbb{P}(\cdot|\mathcal{F}_t)$ is the conditional probability.

The absence of arbitrage (assumed following the definition given by Campbell, 2018) implies the existence of a strictly positive *stochastic discount factor (SDF)* or *pricing kernel* process $\{M_t\}$ (i.e., $M_t \colon \Omega \to \mathbb{R}_{++}$) such that the price S_t of a claim to an arbitrary state-contingent payoff X_T is given by

$$S_t(X_T) = \mathbb{E}_t\left[\frac{M_T}{M_t} X_T\right], \tag{A.2}$$

where $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot|\mathcal{F}_t]$, and we can initialize $M_0 = 1$.

Define the *risk-neutral measure* $\mathbb{P}^*$ with respect to the objective measure $\mathbb{P}$ according to the Radon-Nikodym derivative

$$\left.\frac{d\mathbb{P}^*}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \frac{M_T/M_t}{\mathbb{E}_t[M_T/M_t]}. \tag{A.3}$$

From equation (A.2), the $(T-t)$-period gross risk-free rate is $R^f_{t,T} \equiv 1/S_t(1_T) = 1/\mathbb{E}_t[M_T/M_t]$, where 1_T refers to one unit of the numeraire delivered at T. Using this along with the change of measure in (A.3), rewrite (A.2) as

$$S_t(X_T) = \frac{1}{R^f_{t,T}}\mathbb{E}^*_t[X_T], \tag{A.4}$$

as stated in the text, and where $\mathbb{E}^*_t[\cdot]$ is again the conditional expectation under $\mathbb{P}^*$.

Now, using (A.1) and (A.3), the risk-neutral probability for index state v_j is

$$\mathbb{P}^*_t(V_T = v_j) = \frac{\mathbb{E}_t[M_T/M_t \mid V_T = v_j]}{\mathbb{E}_t[M_T/M_t]}\,\mathbb{P}_t(V_T = v_j). \tag{A.5}$$

The risk-neutral pricing equation (A.4) can then be used to show that the date-t schedule of option prices $\{q_{t,K}\}_K$ reveals the set of risk-neutral probabilities $\{\mathbb{P}^*_t(V_T = v_j)\}_j$, as stated in the text. Assume for notational simplicity that the set of traded option strike prices $\mathcal{K}$ coincides with $\mathcal{V}_T$, and denote $K_j = v_j$ for all j. We can then back out the risk-neutral probabilities of interest from option prices as follows:

$$\mathbb{P}^*_t(V_T = v_j) = R^f_{t,T}\left[\frac{q_{t,K_{j+1}} - q_{t,K_j}}{K_{j+1} - K_j} - \frac{q_{t,K_j} - q_{t,K_{j-1}}}{K_j - K_{j-1}}\right]. \tag{A.6}$$

Augenblick and Lazarus (2022, Appendix A) present a brief derivation of this result, which follows directly from a discrete-state application of the classic result of Breeden and Litzenberger (1978).

Then using the definitions of $\pi_{t,j}$ and $\pi^*_{t,j}$ in equations (1) and (2), respectively, equation (3) in the text then follows immediately from (A.5).

A.2. Proof of Equation (9)

Under the assumption in equation (8), we have

$$\frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1-\widehat{\pi}^*_{t,j})} - \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1-\pi^*_{t,j})}$$

$$= \frac{\pi^*_{t,j} + \epsilon_{t,j}}{\pi^*_{t,j} + \epsilon_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j} - \epsilon_{t,j})} - \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})}$$

$$= \frac{\epsilon_{t,j}\phi_{T-t}}{\left(\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})\right)\left(\epsilon_{t,j}(\phi_{T-t} - 1) + \pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})\right)}.$$

A Taylor expansion of this expression in $\epsilon_{t,j}$ and ϕ_{T-t} around the point $(\|\epsilon_{t,j}\|, \phi_{T-t}) = (0, 1)$, with notation as discussed in Footnote 12, yields

$$\frac{\widehat{\pi}^*_{t,j}}{\widehat{\pi}^*_{t,j} + \phi_{T-t}(1 - \widehat{\pi}^*_{t,j})} - \frac{\pi^*_{t,j}}{\pi^*_{t,j} + \phi_{T-t}(1 - \pi^*_{t,j})}$$

$$= \epsilon_{t,j} + \mathcal{O}\left(\|\epsilon_{t,j}\|^2\right) + \left((2\pi^*_{t,j} - 1)\epsilon_{t,j} + \mathcal{O}\left(\|\epsilon_{t,j}\|^2\right)\right)(\phi_{T-t} - 1) + \mathcal{O}\left((\phi_{T-t} - 1)^2\right)$$

$$= \epsilon_{t,j} + \mathcal{O}\left((\|\epsilon_{t,j}\| + (\phi_{T-t} - 1))^2\right),$$

as stated. □

A.3. Measurement of Risk-Neutral Distribution

I briefly describe the measurement procedure here, and again see Augenblick and Lazarus (2022) for further detail and discussion. In addition to the option prices described in the text, OptionMetrics reports a risk-free zero-coupon yield curve across multiple maturities, as well as the underlying S&P 500 index price. I use the risk-free rate at the relevant horizon as an input in the measurement of risk-neutral beliefs, and I use the index price to observe the ex-post index state for each option expiration date T and assign probability 1 to that state on date T.[26]

I then measure the risk-neutral distribution for returns by applying the following steps to the observed option-price cross-sections, following Malz (2014):

1. Transform the collections of call- and put-price cross-sections (for example, for call options on date t for expiration date T, this set is $\{q_{t,K}\}_{K \in \mathcal{K}}$) into Black–Scholes implied volatilities.

[26] The settlement value for many S&P 500 options in fact reflects the opening (rather than closing) price on the expiration date; for example, the payoff for the traditional monthly S&P 500 option contract expiring on the third Friday of each month depends on the opening S&P index value on that third Friday morning, while the payoff for the more recently introduced end-of-month option contract depends on the closing S&P index value on the last business day of the month. See http://www.cboe.com/SPX for further detail. For my dataset, 441 of the 685 option expiration dates correspond to A.M.-settled options. To obtain the ex-post return state for A.M.-settled options, I hand-collect the option settlement values for these expiration dates from the Chicago Board Options Exchange (CBOE) website, which posts these values.

2. Fit a cubic spline to interpolate a smooth function between the points in the resulting implied-volatility schedule for each trading date–expiration date pair (separately for the call- and put-option values). The spline is *clamped*: its boundary conditions are that the slope of the spline at the minimum and maximum values of the knot points $\{q_{t,K}\}_{K \in \mathcal{K}}$ is equal to 0; further, to extrapolate outside of the range of observed knot points, set the implied volatilities for those unobserved strikes equal to the implied volatility for the closest observed strike (i.e., maintain a slope of 0 for the implied-volatility schedule outside the observed range).

3. Evaluate this spline (separately for calls and puts) at 1,901 strike prices, for S&P index values ranging from 200 to 4,000 (so that the evaluation strike prices are $K = 200, 202, \ldots, 4000$), to obtain a set of implied-volatility values across this fine grid of possible strike prices.[27]

4. Average the separate call- and put-option implied-volatility values from the previous step at each strike for each (t, T) pair, to obtain a single implied-volatility schedule across strikes for each such (t, T) pair. (Given put-call parity, the implied-volatility values for calls and puts should in theory be equal at a given strike; in practice, they tend to differ slightly given market microstructure issues, so using the mean of the two values is a simple way of averaging out the effects of such idiosyncratic noise. This step is the only point of distinction between our procedure and that of Malz, who assumes access to a single implied-volatility schedule and thus does not consider call and put prices separately.)

5. Invert the single resulting smoothed 1,901-point implied-volatility schedule for each (t, T) pair to transform these values back into call prices, and denote this fitted call-price schedule as $\{\hat{q}_{t,K}\}_{K \in \{200,202,\ldots,4000\}}$.

6. Calculate the risk-neutral CDF for the date-T index value at strike price K using $\mathbb{P}^*_t(V_T < K) = 1 + R^f_{t,T}(\hat{q}_{t,K} - \hat{q}_{t,K-2})/2$, following equation (A.6). (The index-value distance between the two adjacent strikes is equal to 2 given that I evaluate the spline at intervals of two index points.)

7. For clarity, temporarily index the set of expiration dates by the subscript i, so that that set is given by $\{T_i\}_i$ (rather than the generic $\{T\}$). Defining $V_{i,j,\max}$ and $V_{i,j,\min}$ to be the date-T_i index values corresponding to the upper and lower bounds, respectively, of the

[27] This set of $\sim$1,900 strike prices is on average about 20 times larger than the set of strikes for which there are prices in the data, as there is a mean of roughly 94 observed values in a typical set $\{q_{t,K}\}_{K \in \mathcal{K}}$ (and similarly for put options).

bin defining index state v_j,[28] I then calculate the risk-neutral probability that state v_j will be realized at date T_i, referred to with slight notational abuse as $\mathbb{P}_t^*(v_j)$, as

$$\mathbb{P}_t^*(s_j) = \mathbb{P}_t^*(V_{T_i} < V_{i,j,\max}) - \mathbb{P}_t^*(V_{T_i} < V_{i,j,\min}),$$

where the CDF values are taken from the previous step using linear interpolation between whichever two strike values $K \in \{200, 202, \ldots, 4000\}$ are nearest to $V_{i,j,\max}$ and $V_{i,j,\min}$, respectively.

Note that I transform the option prices into Black–Scholes implied volatilities simply for purposes of fitting the cubic spline and then transform these implied volatilities back into call prices before calculating risk-neutral beliefs, so this procedure does *not* require the Black–Scholes model to be correct.[29] The clamped cubic spline proposed by Malz (2014), and used in step 2 above, is chosen to ensure that the call-price schedule obtained in step 5 is decreasing and convex with respect to the strike price outside the range of observable strike prices, as required under the restriction of no arbitrage. Violations of these restrictions *inside* the range of observable strikes, as observed infrequently in the data, generate negative implied risk-neutral probabilities; in any case that this occurs, I set the associated risk-neutral probability to 0 and renormalize the remainder of the distribution.

A.4. Details on GMM Estimation and Inference Procedure

I construct risk-price point estimates by horizon, as reported in Figure 1, using two-step GMM. I use the five-day-lagged observation as an instrument, and conduct estimation on the sample counterparts of the moment conditions in equation (12). The price of risk parameter is constrained to be equal for all days within a given weekly horizon to expiration. The first-stage weight matrix is $Z'Z/\mathcal{T}$, where Z is the data matrix for the instruments and $\mathcal{T}$ is the number of observations. The second-stage weight matrix is then clustered by blocks of 8 time-adjacent observations.

This weight-matrix clustering is designed to match the inference procedure for estimating equation (14), which is a block bootstrap with 8-observation (roughly 2-month) blocks. This bootstrap proceeds by re-estimating ϕ_κ on 500 redrawn bootstrap samples, rerunning the regression (14) within each of these samples, and then calculating the distribution of the statistic $\widehat{\beta}^* - \widehat{\beta}$, where $\widehat{\beta}^*$ is the bootstrap estimate for β and $\widehat{\beta}$ is the estimate in the original sample. Denoting by $q^*(\cdot)$ the quantile function of the bootstrap distribution of

[28]That is, formally, $V_{i,j,\min} = R^f_{0_i,T_i} V_{T_0} \exp(s_j - 0.01)$ and $V_{i,j,\max} = R^f_{0_i,T_i} V_{0_i} \exp(s_j + 0.01)$.

[29]I conduct this transformation following Malz (2014), as well as much of the related literature.

$\widehat{\beta}^* - \widehat{\beta}$, I then calculate the 95% confidence interval as $[\widehat{\beta} - q^*(0.975), \widehat{\beta} - q^*(0.025)]$. This follows the standard procedure for handling possible asymmetries in the finite-sample distribution of $\widehat{\beta} - \beta$; see, e.g., Hall (1988), Hansen (2022).

A.5. Tests for Preferences over the Timing of Resolution of Uncertainty

Consider a simple economy with two possible outcomes for the index value at each possible date, again H and L. Consider an option with a date-T payoff of $X_{H,T} = \mathbb{1}\{V_T = H\}$ and date-t price $q_{H,t}$, as well as the complementary low-state option with payoff $X_{L,T} = \mathbb{1}\{V_T = L\} = 1 - X_{H,T}$ and price $q_{L,t}$. Assume for now that $T = 1$. One can construct an early-resolving but late-paying gamble (where the payment horizon is $\overline{T} > 1$) using the following date-0 strategy:

1. Purchase $y/(q_{L,0} + q_{H,0})$ units of the low-state option (which costs $yq_{L,0}/(q_{L,0} + q_{H,0}) = y\pi_0^*$), for a value y to be determined below.
2. Purchase a forward contract to invest the proceeds of the date-1 option payoff from the previous step in the forward rate from $t = 1$ to $\overline{T}$, denoted $f_{0,1,\overline{T}}$, conditional on the date-1 state being L. Set $y = 1/f_{0,1,\overline{T}}$ in the previous step.

It can be seen that this strategy pays off 1 at date T if state L is revealed to be realized at date 1, and 0 otherwise. While the forward contract in step 2 is not directly observable in the data, its (average) price can be inferred using ex-post realizations of the risk-free rate $R^f_{1,\overline{T}}$ in state L, using an uncovered interest parity–like argument. This thus in theory allows for a test of preferences over the timing of resolution of uncertainty.

References

AÏT-SAHALIA, Y. AND A. W. LO (1998): "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, 53, 499–547.

——— (2000): "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics*, 94, 9–51.

ANDRIES, M., T. M. EISENBACH, AND M. C. SCHMALZ (2019): "Horizon-Dependent Risk Aversion and the Timing and Pricing of Uncertainty," *Federal Reserve Bank of New York Staff Report No. 703*.

AUGENBLICK, N. AND E. LAZARUS (2022): "A New Test of Excess Movement in Asset Prices," *Working Paper*.

AUGENBLICK, N. AND M. RABIN (2021): "Belief Movement, Uncertainty Reduction, and Rational Updating," *Quarterly Journal of Economics*, 136, 933–985.

BANSAL, R., S. MILLER, D. SONG, AND A. YARON (2021): "The Term Structure of Equity Risk Premia," *Journal of Financial Economics*, 142, 1209–1228.

BANSAL, R. AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

BARBERIS, N. AND M. HUANG (2008): "The Loss Aversion/Narrow Framing Approach to the Equity Premium Puzzle," in *Handbook of the Equity Risk Premium*, ed. by R. Mehra, Amsterdam: Elsevier, chap. 6, 199–236.

BENARTZI, S. AND R. H. THALER (1995): "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics*, 110, 73–92.

VAN BINSBERGEN, J. H., M. BRANDT, AND R. KOIJEN (2012): "On the Timing and Pricing of Dividends," *American Economic Review*, 102, 1596–1618.

BINSBERGEN, J. V., W. HUESKES, R. KOIJEN, AND E. VRUGT (2013): "Equity Yields," *Journal of Financial Economics*, 110, 503–519.

BINSBERGEN, J. V. AND R. KOIJEN (2017): "The Term Structure of Returns: Facts and Theory," *Journal of Financial Economics*, 124, 1–21.

BLACK, F. AND M. SCHOLES (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637.

BLISS, R. R. AND N. PANIGIRTZOGLOU (2004): "Option-Implied Risk Aversion Estimates," *Journal of Finance*, 59, 407–446.

BOGUTH, O., M. CARLSON, A. FISHER, AND M. SIMUTIN (2011): "Conditional Risk and Performance Evaluation: Volatility Timing, Overconditioning, and New Estimates of Momentum Alphas," *Journal of Financial Economics*, 102, 363–389.

——— (2022): "The Term Structure of Equity Risk Premia: Levered Noise and New Estimates," *Forthcoming, Review of Finance*.

BOROVIČKA, J., L. P. HANSEN, AND J. A. SCHEINKMAN (2016): "Misspecified Recovery," *Journal of Finance*, 71, 2493–2544.

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

CABALLERO, R. J., E. FARHI, AND P.-O. GOURINCHAS (2017): "Rents, Technical Change, and Risk Premia: Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares," *American Economic Review: Papers & Proceedings*, 107, 614–620.

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

CAMPBELL, J. Y. AND T. VUOLTEENAHO (2004): "Bad Beta, Good Beta," *American Economic Review*, 94, 1249–1275.

CHABI-YO, F., R. GARCIA, AND E. RENAULT (2008): "State Dependence Can Explain the Risk Aversion Puzzle," *Review of Financial Studies*, 21, 973–1011.

DEW-BECKER, I., S. GIGLIO, A. LE, AND M. RODRIGUEZ (2017): "The Price of Variance Risk," *Journal of Financial Economics*, 123, 225–250.

EISENBACH, T. M. AND M. C. SCHMALZ (2016): "Anxiety in the Face of Risk," *Journal of Financial Economics*, 121, 414–426.

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

GIGLIO, S., M. MAGGIORI, AND J. STROEBEL (2015): "Very Long-Run Discount Rates," *Quarterly Journal of Economics*, 130, 1–53.

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

HALL, P. (1988): "Theoretical Comparison of Bootstrap Confidence Intervals," *Annals of Statistics*, 16, 927–953.

HALL, P. AND J. L. HOROWITZ (1996): "Bootstrap Critical Values for Tests Based on Generalized-Method-of-Moments Estimators," *Econometrica*, 64, 891–916.

HANSEN, B. (2022): *Econometrics*, Princeton, NJ: Princeton University Press.

HANSEN, L. P. (1982): "Large Sample Properties of Generalised Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

——— (1985): "A Method for Calculating Bounds on the Asymptotic Covariance Matrices of Generalized Method of Moments Estimators," *Journal of Econometrics*, 30, 203–238.

HANSEN, L. P. AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.

HANSEN, L. P. AND K. J. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1286.

HASLER, M. AND R. MARFÈ (2016): "Disaster Recovery and the Term Structure of Dividend Strips," *Journal of Financial Economics*, 122, 116–134.

HAYASHI, F. AND C. SIMS (1983): "Nearly Efficient Estimation of Time Series Models with Predetermined, but not Exogenous, Instruments," *Econometrica*, 51, 783–798.

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

KREPS, D. M. AND E. L. PORTEUS (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185.

LAZARUS, E., D. J. LEWIS, AND J. H. STOCK (2021): "The Size-Power Tradeoff in HAR Inference," *Econometrica*, 89, 2497–2516.

LETTAU, M. AND J. A. WACHTER (2007): "Why is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium," *Journal of Finance*, 62, 55–92.

LICHTENSTEIN, S., B. FISCHHOFF, AND L. D. PHILLIPS (1977): "Calibration of Probabilities: The State of the Art," in *Decision Making and Change in Human Affairs: Proceedings of the Fifth Research Conference on Subjective Probability, Utility, and Decision Making*, ed. by H. Jungermann and G. De Zeeuw, Dordrecht: Springer Netherlands, 275–324.

LINN, M., S. SHIVE, AND T. SHUMWAY (2018): "Pricing Kernel Monotonicity and Conditional Information," *Review of Financial Studies*, 31, 493–531.

LOEWENSTEIN, G. (1996): "Out of Control: Visceral Influence on Behavior," *Organizational Behavior and Human Decision Processes*, 65, 272–292.

MALZ, A. M. (2014): "A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions," *Federal Reserve Bank of New York Staff Report No. 677*.

MARTIN, I. (2017): "What Is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132, 367–433.

MEHRA, R. AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.

PLAGBORG-MØLLER, M. (2016): "Essays in Macroeconometrics," Ph.D. thesis, Harvard University.

POHL, W., K. SCHMEDDERS, AND O. WILMS (2018): "Higher Order Effects in Asset Pricing Models with Long-Run Risks," *Journal of Finance*, 73, 1061–1111.

RABIN, M. AND G. WEIZSÄCKER (2009): "Narrow Bracketing and Dominated Choices," *American Economic Review*, 99, 1508–1543.

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

ROSS, S. (2015): "The Recovery Theorem," *Journal of Finance*, 70, 615–648.

SONG, Y. (2018): "Dealer Funding Costs: Implications for the Term Structure of Dividend Risk Premia," *University of Washington Mimeo.*

SUN, Y. AND M. S. KIM (2012): "Simple and Powerful GMM Over-Identification Tests with Accurate Size," *Journal of Econometrics*, 166, 267–281.

TVERSKY, A. AND D. KAHNEMAN (1981): "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 453–458.

WACHTER, J. A. (2002): "Comment on: Are Behavioral Asset-Pricing Models Structural?" *Journal of Monetary Economics*, 49, 229–233.

WEBER, M. (2018): "Cash Flow Duration and the Term Structure of Equity Returns," *Journal of Financial Economics*, 128, 486–503.