# High Volatility is the Key: A Joint Treatment of Asset Pricing Puzzles*

Errikos Melissinos†

April 2026

Earlier Version of Article

**Abstract**

Macrofinance faces a fundamental tension in that aggregate consumption volatility is low, yet risk premia are high. However, increasing evidence suggests that regular households are not marginal investors. This article models the marginal investor as a consumer with high, time-varying consumption volatility, motivated by the intermediary asset pricing perspective. This approach yields a joint explanation of five puzzles — the equity premium, the risk-free rate, the bond premium, stock return predictability, and excess volatility. Such stochastic volatility appears to be a necessary ingredient for jointly explaining these puzzles within any standard single-state-variable consumption-based model.

**JEL:** C65, E43, E44, G12

**Keywords:** term premia, stochastic volatility, intermediary asset pricing, high consumption volatility, equity premium, return predictability, excess volatility, habit, long-run risk

*For helpful comments I thank Mariia Bondar, Julian Detemple, Fabio Girardi, Refet Gürkaynak, Michael Haliassos, Florian Heider, Alexander Hillert, Ruggero Jappeli, Martin Kliem, Holger Kraft, Nora Lamersdorf, Mortiz Lenel, Martin Lettau, Christoph Meinedring, Alexander Meyer-Gohde, Andrea Modena, Emanuel Mönch, Dmitry Mukhin, Christian Mücke, Han Özsoylev, Loriana Pelizzon, Christian Schlag, Paul Schneider, Nicolas Syrichas, Farah Tohme, Henning Weber, Volker Wieland, the participants of the SAFE Brown Bag Seminar and the participants of the first Bonn-Frankfurt-Mannheim PhD Conference.

†Özyegin University
website: https://errikos-melissinos.com

# 1 Introduction

A long-standing challenge in asset pricing is explaining several empirical regularities that standard representative-agent models fail to reproduce. These models typically assume a representative consumer, yet various strands of literature independently suggest that the marginal investor differs from the representative household. Regular households have limited stock market participation (Haliassos and Bertaut 1995; Campbell 2006); supply-and-demand-driven approaches to asset pricing show that agents such as "arbitrageurs" play a special role in price formation (Vayanos and Vila 2021); and the intermediary asset pricing approach identifies financial intermediaries as the relevant marginal investors (He and Krishnamurthy 2013). Building on the intermediary perspective, this article models the marginal investor as a consumer with high, time-varying consumption volatility, and shows that this single feature jointly accounts for several prominent asset pricing puzzles.

Asset pricing puzzles arise when standard macroeconomic models fail to replicate observed patterns in central financial variables — such as stock returns, the risk-free rate, and bond yields — either qualitatively or in magnitude.[1] The equity premium puzzle, introduced by Mehra and Prescott (1985), is the finding that excess stock market returns cannot be explained by the standard consumption-based model with moderate risk aversion when calibrated to aggregate consumption data. Attempting to resolve this puzzle by raising risk aversion leads to a counterfactually volatile risk-free rate, a further inconsistency known as the risk-free rate puzzle (Weil 1989).

Furthermore, stock returns demonstrate high variability and predictability (Campbell and Shiller 1988; Shiller 1981), giving rise respectively to the excess volatility puzzle and the return predictability puzzle. Apart from stocks, long-term bonds also exhibit positive excess returns, which are large and time-varying, while most standard models predict excess returns for bonds that are negative and small in absolute value. This is the bond premium puzzle (Backus *et al.* 1989).[2] These bond features are also closely related to the term structure of interest rates, as positive bond premia are associated with an upward sloping yield curve. Originally, the puzzle was introduced for nominal bonds. For real (inflation-adjusted) bonds, evidence is mixed: the real premium appears negative in the UK (Piazzesi and Schneider 2006), while in the USA most evidence points to a significant, mostly positive and time-varying real term premium (Abrahams *et al.* 2016; d' Amico *et al.*

[1] A good summary of these puzzles is provided in Gabaix (2012).

[2] The bond premium puzzle in the literature often refers to term premia. These are excess returns of long-term bonds over short-term bonds for an investment horizon which is equal to the maturity of the long-term bond, and they are defined precisely later in the article.

2018; Pflueger and Viceira 2016). Hence, generating a positive and time-varying real term premium consistent with the US evidence is a key objective of this article.

The intermediary asset pricing approach has shown promise in explaining some of these puzzles. For example, in foundational articles such as He and Krishnamurthy (2013) and Basak and Cuoco (1998), a significant equity premium is generated in certain states of the world.[3] Adrian *et al.* (2014), He *et al.* (2017), and Haddad and Muir (2021) further demonstrate the empirical relevance of intermediary asset pricing in explaining the cross-section of asset returns. However, despite a large literature on intermediary asset pricing, a joint treatment of the puzzles mentioned above is missing.

This article uses an exogenous consumption-based process, where "consumption" can also be interpreted as the intermediary's performance, with characteristics chosen to fit the intermediaries in He *et al.* (2017). In particular, consumption volatility is taken to be large, time-varying, and negatively correlated with stochastic consumption growth. This follows the intuition that intermediaries' risk exposure increases during distress and is considerably larger compared to regular households. In He *et al.* (2017), intermediaries are distressed when the intermediary capital ratio is low. Likewise, in the current article, distress coincides with high consumption volatility and negative consumption shocks. An alternative intermediary perspective is offered by Adrian and Shin (2014), who document procyclical leverage: in their framework, intermediary balance sheets expand in good times, so that high leverage coincides with strong intermediary performance rather than distress. This would correspond to a positive correlation between consumption volatility and consumption in the present model. The two perspectives make opposite sign predictions for this correlation, and the model's ability to match empirical moments is sensitive to which restriction is imposed, as discussed further in the results. By using an exogenous consumption process, the model abstracts from the specific intermediation frictions. Indeed, the model is in a sense more general than the intermediary asset pricing approach, as it applies to any situation where the marginal investor has similarly high and time-varying consumption volatility.[4] For instance, Mitra and Xu (2024) generate bond premia in a setting where households have high consumption volatility by not being able to share idiosyncratic risk.

[3]Basak and Cuoco (1998) is primarily a limited stock market participation model, where the marginal investors are unconstrained households rather than financial intermediaries in the strict sense. Nevertheless, the unconstrained agents in that paper are reminiscent of intermediaries.

[4]For example, Ait-Sahalia *et al.* (2004) provide evidence that the average consumption volatility of wealthy individuals may be considerably higher than standard measures of aggregate consumption volatility.

Nevertheless, the intermediary perspective is the main motivation for the model, and it is used to provide intuition for the results. Thus, the article also contributes by demonstrating how the intermediary perspective fits within a simple consumption-based framework.

The exogenous state variable that drives consumption volatility also drives the consumption drift. In particular, in a downturn there is an expectation of a rebound, and distressed intermediaries may require higher expected returns to compensate for higher risk, leading to a higher consumption drift. The risk-free interest rate is positively related to consumption drift through the consumption smoothing motive and negatively related to consumption volatility through the precautionary savings motive. Given that both these quantities are increasing in the state variable, the effect of consumption drift tends to dominate the precautionary savings effect over the range of the state variable most relevant for risk premia, so that the risk-free rate is on net positively related to the state variable.[5] These features generate a high and time-varying equity premium, a low and time-varying risk-free rate by construction, and a high and time-varying bond premium. The bond premium is positive because the risk-free rate is mostly negatively correlated with the consumption process of the intermediary, which makes bonds perform well (badly) in good (bad) times for the intermediaries. Beyond the level of equity and bond premia, the model also generates significant return predictability and volatility of excess stock returns, while also capturing other features of the term structure of interest rates.

This article also argues that high consumption volatility is an essential ingredient common to many successful explanations of these puzzles. Various approaches, including limited stock market participation (Basak and Cuoco 1998), intermediary asset pricing (He and Krishnamurthy 2013 and more recently relating to the bond premium puzzle Kekre *et al.* 2022; Schneider 2022), and supply-and-demand driven models (Vayanos and Vila 2021), all converge on the idea that the marginal investor faces high, time-varying volatility in consumption or wealth. The current article distills these common features into a more abstract consumption-based model, and shows that such a model jointly explains a series of asset pricing puzzles.

Notably, addressing the bond premium puzzle is particularly challenging. The difficulty lies in the fact that real bond prices move inversely with interest rates. Since interest rates are typically pro-cyclical, real bonds should theoretically hedge

[5]More precisely, the calibrated risk-free rate exhibits a U-shaped profile with a minimum slightly below the steady state; see Section 4. This net positive relationship is necessary for generating a positive bond premium.

consumption risk and command a negative risk premium instead of a positive one. As shown in the main part of this article, even recursive utility models fail to account for a positive real bond premium, despite their success in explaining the equity premium puzzle (Bansal and Yaron 2004). The rare disasters literature (Rietz, 1988; Barro, 2006) has also shown promise in explaining the equity premium and other asset pricing puzzles. Gabaix (2012) even generates a positive and sizable nominal bond premium by assuming that inflation is on average high during consumption disasters. However, the underlying real bond premium in this framework remains negative, so the challenge of explaining a positive real term premium persists. The reason is that the forces that generate the high equity premium can increase the absolute value of the bond premium, but they cannot by themselves make it positive rather than negative. Beyond the high consumption volatility approach, there are two main alternatives.[6] The first relies on an inflation channel, which by definition does not apply to real bonds. The second alternative, which is able to generate positive real term premia, relies on highly time-varying effective risk aversion, such as in habit-based frameworks (Campbell and Cochrane 1999; Wachter 2006).[7] However, these models have been criticized with respect to the equity premium puzzle, because they also require an implausibly high effective risk aversion, as was originally pointed out when the equity premium puzzle was introduced (Mehra 2007).

In the following, Section 2 introduces the model, Section 3 discusses the calibration, and Section 4 presents the core findings. These include the predictions of the model related to the equity premium, the risk-free rate, stock market predictability, excess volatility, and the bond/term premium. Section 5 then discusses the special case of recursive preferences, demonstrating that models with more standard preference specifications and a single state variable cannot jointly match the empirical moments, in particular the positive real term premium. Section 6 discusses the plausibility of high marginal-investor consumption volatility and relates the model to indirect empirical evidence from intermediary data. Section 7 concludes.

[6] A more extensive review of this literature is provided in Melissinos (2023b), which is an earlier version of this article.

[7] It is also possible to account for the puzzle by using explicit time-varying risk aversion (Lettau and Wachter 2011) or other non-standard preferences (Zhao 2020; Ellison and Tischbirek 2021).

# 2 Model Setup

## 2.1 One Marginal Investor

Similar to He *et al.* (2017), the model focuses on a single agent who is the marginal investor, abstracts from non-marginal investors, and takes the consumption process of this agent to be exogenous.

Theoretically, this could be the only agent in the economy, but it is more realistic to think that other agents also exist but do not participate in the same asset market due to frictions. While it makes more sense to think of the marginal investors as some kind of financial intermediary, it is also possible to imagine that a small group of households simultaneously participate in financial markets and have a high consumption volatility. For example, Ait-Sahalia *et al.* (2004) provide evidence that high net-worth households have a much higher consumption volatility than regular households.

## 2.2 Utility

The lifetime utility of the agent is given by:

$$U_0 = \mathrm{E}_0 \int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma} - 1}{1-\gamma} \mathrm{d}t, \tag{1}$$

where $C_t$ is the consumption flow at time $t$, $\gamma$ is the coefficient of relative risk aversion, and $\rho$ is the time preference rate. In the model, time is continuous, utility is time-separable, and the utility flow exhibits constant relative risk aversion (CRRA).

## 2.3 Consumption and State Processes

The consumption flow $C_t$ is exogenous and follows a stochastic process:

$$\mathrm{d}\log(C_t) = \mathrm{d}c_t = \mu_{ct}\mathrm{d}t + \sigma_{ct}\mathrm{d}W_{ct} \tag{2}$$

where $c_t$ is log consumption, $x_t$ is the state variable, $W_{ct}$ is a Wiener process associated with consumption, $\mu_{ct}$ is the consumption drift, and $\sigma_{ct}$ is the consumption diffusion. The specific functional forms of $\mu_{ct}$ and $\sigma_{ct}$ are given in the Appendix; they depend on parameters whose calibration is described in Section 3. As a result of the calibration, both $\mu_{ct}$ and $\sigma_{ct}$ turn out to be increasing in the state variable $x_t$. The most important feature of the consumption process is its high time-varying

volatility, which as shown later contributes to a high time-varying equity premium and a large bond premium in absolute value.

There is only one state variable $x_t$, and it follows a mean-reverting process:

$$\begin{aligned} \mathrm{d}x_t &= \log(\phi)x_t\mathrm{d}t + \sigma_{xt}\mathrm{d}W_{xt} \\ \mathrm{d}W_{xt}\mathrm{d}W_{ct} &= \rho_{cxt}\mathrm{d}t \end{aligned} \tag{3}$$

where $\sigma_{xt}$ is the state variable diffusion, $W_{xt}$ is a Wiener process associated with the state variable, and $\rho_{cxt}$ is the correlation between the noise processes for consumption and the state variable. The drift term shows that the steady state is at $x_t = 0$, and the process always reverts to the steady state given that $0 < \phi < 1$.[8]

## 2.4 Stochastic Discount Factor

Based on the utility function, the stochastic discount factor (SDF) is given by:

$$\Lambda_t = e^{-\rho t}C_t^{-\gamma} \tag{4}$$

Then based on the consumption process, and by applying Itô's lemma, the dynamics of the SDF can be derived:[9]

$$\frac{\mathrm{d}\Lambda_t}{\Lambda_t} = \Big( -\rho - \gamma\mu_{ct} + \frac{\gamma^2}{2}\sigma_{ct}^2\Big)\mathrm{d}t - \gamma\sigma_{ct}\mathrm{d}W_{ct} \tag{5}$$

The instantaneous risk-free rate then follows as:

$$r(x_t)\mathrm{d}t = -\mathrm{E}_t\left[\frac{\mathrm{d}\Lambda_t}{\Lambda_t}\right] = \Big(\rho + \gamma\mu_{ct} - \frac{\gamma^2}{2}\sigma_{ct}^2\Big)\mathrm{d}t \tag{6}$$

## 2.5 Zero-Coupon Bond Pricing

The SDF derived above can be used to price any asset via the standard pricing equation. In particular, a zero-coupon bond that pays \$1 at maturity $T = t + m$ must satisfy:

$$\mathrm{E}\big[\mathrm{d}(\Lambda_t B_t(x, m))\big] = 0, \quad B(x, 0) = 1 \tag{7}$$

where $B(x, m)$ is the price of a zero-coupon bond with a remaining maturity of $m$ at time $t$. The price of the bond is a function of the state variable $x_t$, and maturity

[8] This is similar to an AR(1) where $x_{t+1} = \phi x_t + \epsilon_t$.

[9] The pricing follows the notation and the approach of Cochrane (2009).

$m$. Using Itô's lemma on the bond price yields the following:

$$\mathrm{d}B_t = \left(\log(\phi)x_t B_x - B_m + \frac{\sigma_{xt}^2}{2}B_{xx}\right)\mathrm{d}t + \sigma_{xt}B_x \mathrm{d}W_{xt}, \quad B(x,0) = 1 \tag{8}$$

where subscripts on $B$ denote partial derivatives with respect to the corresponding variable.[10] Using the process for the bond price, the pricing equation gives rise to a partial differential equation (PDE) (time subscripts and the arguments of $B$ are omitted for brevity):

$$\begin{aligned} &\mathrm{E}\big[\mathrm{d}(\Lambda B)\big] = 0 \Rightarrow \mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}B + \mathrm{d}B + \frac{\mathrm{d}\Lambda}{\Lambda}\mathrm{d}B\right] = 0 \\ &\Rightarrow -B_m - r(x)B + \big(\log(\phi)x - \rho_{cx}\gamma\sigma_c\sigma_x\big)B_x + \frac{\sigma_x^2}{2}B_{xx} = 0 \end{aligned} \tag{9}$$

where $B_m$ is the derivative of the bond price with respect to maturity, $B_x$ is the derivative of the bond price with respect to the state variable, and $B_{xx}$ is the second derivative of the bond price with respect to the state variable. This differential equation is solved numerically via the Feynman-Kac formula (see Appendix A). Based on the expressions above, the instantaneous expected excess return of the bond before maturity can be derived as:

$$\begin{aligned} \mathrm{E}\left[\frac{\mathrm{d}B}{B}\right] + \mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}\right] = \underbrace{\mathrm{E}\left[\frac{\mathrm{d}B}{B}\right] - r\mathrm{d}t}_{\equiv\text{Expected Excess Return}} &= -\mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}\frac{\mathrm{d}B}{B}\right] \\ \Rightarrow -\frac{B_m}{B} + \log(\phi)x\frac{B_x}{B} + \frac{\sigma_x^2}{2}\frac{B_{xx}}{B} - r &= \rho_{cx}\gamma\sigma_c\sigma_x\frac{B_x}{B} \end{aligned} \tag{10}$$

The term premium is defined as:[11]

$$TP(x_t, m) = -\frac{1}{m}\log\left[\frac{B(x_t,m)}{\mathrm{E}_t\left[\exp\left\{-\int_0^m r(x_{t+s})\mathrm{d}s\right\}\right]}\right] \tag{11}$$

Intuitively, this expresses the expected annualized excess return of the bond over a period equal to its remaining maturity. The *risk-neutral yield* of the bond is correspondingly defined as the yield that would prevail in the absence of any risk

[10] For simplicity the dependence on time is not explicitly shown, but maturity is connected to time through $\partial/\partial t = -\partial/\partial m$.

[11] Based on the Feynman-Kac formula, the equation can also be written like: $TP(x_t, m) = -\frac{1}{m}\log\left[\frac{\mathrm{E}_t^*\left[\exp\left\{-\int_0^m r(x_{t+s})\mathrm{d}s\right\}\right]}{\mathrm{E}_t\left[\exp\left\{-\int_0^m r(x_{t+s})\mathrm{d}s\right\}\right]}\right]$, where the numerator uses the risk-neutral measure.

premium:

$$RNY(x_t, m) = -\frac{1}{m} \log \left[ \mathrm{E}_t \left[ \exp \left\{ - \int_0^m r(x_{t+s}) \mathrm{d}s \right\} \right] \right] \tag{12}$$

so that the total yield $y(x_t, m) = -\frac{1}{m} \log[B(x_t, m)]$ decomposes as $y(x_t, m) = RNY(x_t, m) + TP(x_t, m)$. Equation (10) is informative because it directly measures the bond's expected excess return. In addition, based on Equation (9) $\rho_{cxt}\gamma\sigma_{ct}\sigma_{xt}$ can be compared to $\log(\phi)x_t$. The former is the *term premium* component, as it drives the term premium, while the latter term is the *expectation* component, as it drives deviation of yields from the short-term rate that is due to expected changes in the short-term in the future. If the term premium component is negligible in size compared to the expectation component, then the term premium is zero for practical purposes. For example, if the correlation between the noise processes of consumption and the state variable ($\rho_{cxt}$) is zero, then the term premium is zero and the expectations hypothesis holds.[12] The formula also shows that there is no easy way to generate a sizeable term premium in this type of model, other than having high consumption volatility or high risk aversion. An earlier version of this article (Melissinos (2023b)) checks several model variations, confirming that standard models fail to generate a sizeable term premium; habit-based models (Campbell and Cochrane 1999; Wachter 2006) are an exception, as noted earlier. In these models, the sign of the bond premium comes down to a simple parameter choice that regulates the strength of the intertemporal substitution effect, as opposed to the precautionary savings effect. In Section 5, models with recursive preferences are also shown to fail unless consumption volatility follows the patterns described in this article.

## 2.6 Equity Pricing

Having defined bond prices, a similar approach is applied to stocks. First, stocks are assumed to pay a dividend that follows the process:

$$\begin{aligned} \frac{\mathrm{d}D_t}{D_t} &= \mu_D \mathrm{d}t + \sigma_D \mathrm{d}W_{Dt} \\ \mathrm{d}W_{Dt}\mathrm{d}W_{ct} = \rho_{cD}\mathrm{d}t, &\quad \mathrm{d}W_{Dt}\mathrm{d}W_{xt} = \rho_{xD}\mathrm{d}t \end{aligned} \tag{13}$$

[12] Indeed, it holds in its strongest form, as the term premium is exactly zero. In the literature, the expectations hypothesis usually refers to whether the term premium is changing over time, and not whether it is exactly zero.

where $D_t$ is the dividend flow at time $t$, $\mu_D$ is the drift of the dividend process, $\sigma_D$ is the dividend diffusion, $W_{Dt}$ is a Wiener process associated with the dividend process, $\rho_{cD}$ is the correlation between the noises processes for consumption and the dividend, and $\rho_{xD}$ is the correlation between the noises processes for the state variable and the dividend. For simplicity, $\mu_D$ and $\sigma_D$ are assumed to be constant.[13] The stock price is built by first pricing individual dividend strips and then aggregating across all maturities. A dividend strip with maturity $m$ is a security that pays the dividend $D_{t+m}$ at time $t+m$; let $\hat{S}(x_t, m)$ denote its price at time $t$.

It is convenient to work with the strip price-dividend ratio $\hat{s}(x_t, m) = \hat{S}(x_t, m)/D_t$, normalized by the current dividend. This normalization yields a clean terminal condition. At maturity $m = 0$ the strip pays the current dividend, so $\hat{s}(x, 0) = 1$, which uniquely pins down the solution.[14] By Itô's lemma, $\hat{s}$ follows the process:

$$\mathrm{d}\hat{s}_t = \left(\log(\phi) x_t \hat{s}_x - \hat{s}_m + \frac{\sigma_{xt}^2}{2}\hat{s}_{xx}\right)\mathrm{d}t + \sigma_{xt}\hat{s}_x \mathrm{d}W_{xt} \tag{14}$$

The strip price-dividend ratio satisfies a pricing equation analogous to the bond case (time subscripts and the arguments of $\hat{s}$ are omitted for brevity):[15]

$$\begin{aligned}
&\mathrm{E}\big[\mathrm{d}(\Lambda\hat{S})\big] = 0 \\
\Rightarrow &\mathrm{E}\big[\mathrm{d}(\Lambda \underbrace{(\hat{s}D)}_{\substack{\text{price of}\\ \text{dividend}\\ \text{strip}}})\big] = 0 \Rightarrow \mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}\hat{s} + \mathrm{d}\hat{s} + \frac{\mathrm{d}D}{D}\hat{s} + \frac{\mathrm{d}\Lambda}{\Lambda}\mathrm{d}\hat{s} + \frac{\mathrm{d}\Lambda}{\Lambda}\frac{\mathrm{d}D}{D}\hat{s} + \mathrm{d}\hat{s}\frac{\mathrm{d}D}{D}\right] = 0 \\
\Rightarrow &-\hat{s}_m - \underbrace{\left(r(x) - \mu_D + \rho_{cD}\gamma\sigma_c\sigma_D\right)}_{\hat{r}(x)}\hat{s} + \left(\log(\phi)x - \rho_{cx}\gamma\sigma_c\sigma_x + \rho_{xD}\sigma_x\sigma_D\right)\hat{s}_x + \frac{\sigma_x^2}{2}\hat{s}_{xx} = 0 \\
&\hat{s}(x, 0) = 1
\end{aligned} \tag{15}$$

where $D$ is the dividend paid at time $t+m$. The differential equation can be solved in the same way as the bond pricing equation. The risk premium of the dividend

[13] The model can still be solved in a similar way if they depend on the state variable $x_t$.

[14] Trying to determine the price-dividend ratio for the stock directly gives an equation that does not have such a clear terminal equation. Solving a differential equation of the price-dividend ratio directly would require specifying two conditions Chen *et al.* (2010); the approach followed here avoids this.

[15] This is slightly abusing notation. To be fully precise, an extra time subscript is required, i.e. $\hat{s}_t(x_t, m)$, in order to specify the timing of the dividend. In this notation, maturity occurs at time $t+m$. However, for brevity, maturity is assumed to be given and the time subscript is omitted.

strip can also be derived from the pricing equation:[16]

$$\mathrm{E}\left[\frac{\mathrm{d}(\hat{s}D)}{\hat{s}D}\right]+\mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}\right]=\underbrace{\mathrm{E}\left[\frac{\mathrm{d}(\hat{s}D)}{\hat{s}D}\right]-r(x)\mathrm{d}t}_{\equiv\text{Expected Excess Return}}=-\mathrm{E}\left[\frac{\mathrm{d}\Lambda}{\Lambda}\frac{\mathrm{d}\hat{s}}{\hat{s}}+\frac{\mathrm{d}\Lambda}{\Lambda}\frac{\mathrm{d}D}{D}\right] \tag{16}$$

$$-\frac{\hat{s}_m}{\hat{s}}+\log(\phi)x\frac{\hat{s}_x}{\hat{s}}+\frac{\sigma_x^2}{2}\frac{\hat{s}_{xx}}{\hat{s}}+\mu_D+\rho_{xD}\sigma_D\sigma_x\frac{\hat{s}_x}{\hat{s}}-r=\rho_{cx}\gamma\sigma_c\sigma_x\frac{\hat{s}_x}{\hat{s}}+\rho_{cD}\gamma\sigma_c\sigma_D$$

The stock price-dividend ratio is obtained by integrating the strip price-dividend ratio across all maturities:[17]

$$\frac{S(x_t)}{D(x_t)}=s(x_t)=\int_0^\infty\hat{s}(x_t,m)\mathrm{d}m \tag{17}$$

where the stock's dividend flow matches the dividends paid by the dividend strips.[18] The expected excess return of the stock follows correspondingly as a weighted average of the strip excess returns across maturities:[19]

$$\mathrm{E}\left[\frac{\mathrm{d}(sD)}{sD}\right]-r\,\mathrm{d}t=\frac{1}{s(x_t)}\int_0^\infty\left(\mathrm{E}\left[\frac{\mathrm{d}(\hat{s}D)}{\hat{s}D}\right]-r\,\mathrm{d}t\right)\hat{s}(x_t,m)\,\mathrm{d}m \tag{18}$$

# 3 Calibration

Having laid out the model, this section describes how its parameters are calibrated. Some parameters are set directly from the data or calibrated independently of the targeting procedure, while others are calibrated by targeting empirical moments, with the targeting procedure described in detail in Appendix D.1. Time is measured in years throughout.

**Basic parameters.** The relative risk aversion of the marginal investor is set to $\gamma = 2$, a moderate value by asset-pricing standards. Values of $\gamma$ between 5 and 10 are common in the relevant literature. The fact that the model generates

[16]The risk premium is used equivalently as the expected excess return.

[17]A similar approach is used in Wachter (2006).

[18]In the case of a dividend strip, a lump sum amount is paid at a specific point in time, while the stock pays a dividend flow. For example, in the case of a dividend strip, if $D_0 = 1$ the dividend strip pays \$1 at time $t = 0$, and, in the case of a stock, if $D_t = 1$ for all $t \in (0, 1)$, then the stock pays a constant flow of dividends summing up to \$1 over the period from $t = 0$ to $t = 1$.

[19]Alternatively it can also be found directly through the pricing equation of the stock. Following Cochrane (2009) the pricing equation for a security with a dividend flow is given by $\mathrm{E}\big[\mathrm{d}(\Lambda S)\big] + \Lambda D\mathrm{d}t = 0$, where $S$ is the price of the security, and $D$ is the dividend flow. This equation can be rearranged to give the expected excess return.

non-trivial equity and bond term premia at $\gamma = 2$ underscores that the mechanism operates through the state-dependence of consumption volatility, not through an elevated curvature of the utility function. The steady-state level of the state variable is normalized to $\bar{x} = 0$ without loss of generality, and its diffusion is set to $\sigma_x = 0.2$, a value chosen so that the state variable typically takes values within the interval $(-1, +1)$. The model is defined on the entire real line; $x = \pm 1$ serve as the reference points for the consumption-volatility and risk-free rate schedules in the calibration. The dividend process has drift $\mu_D = 2.89\%$ and diffusion $\sigma_D = 1.60\%$.[20] The correlations between dividend shocks and both consumption shocks and state-variable shocks are set to zero ($\rho_{cD} = \rho_{xD} = 0$), for two reasons: simplicity, and to allow the consumption-volatility channel to operate in isolation without being confounded by dividend-shock transmission.

The remaining parameters—the state-variable persistence $\phi$, the consumption–state correlation $\rho_{cx}$, the consumption-volatility schedule at the two boundaries ($\sigma_c^{\text{lo}}$, $\sigma_c^{\text{hi}}$), the risk-free rate at the two boundaries ($r^{\text{lo}}$, $r^{\text{hi}}$), and the time-preference rate $\rho$—are calibrated by targeting empirical moments. For a given set of target moments, parameters are chosen to minimize a sum of squared deviations between simulated and empirical moments, using the Nelder–Mead algorithm (Mogensen and Riseth, 2018) applied to bounded parameters. A systematic sweep of 24 moment combinations is run to characterise which moments the model can jointly match and to assess the robustness of the qualitative results; the full sweep design, targeting procedure, and results are described in Appendix D.

The baseline used throughout the article targets two empirical moments: the mean P/D ratio, drawn from Shiller (2015) over the period 1990–2024, and the TIPS 10–5 yr yield spread, derived from the zero-coupon real yield curve of Gürkaynak *et al.* (2010) over the period 2004–2024. The remaining moments are not targeted. Table 1 reports the calibrated parameter vector. As shown in Appendix D, targeting alternative moment combinations yields qualitatively similar results across many runs.

One structural restriction is imposed separately from the targeting procedure. The sign of $\rho_{cx}$ is required to be negative. In the model, $\rho_{cx} < 0$ means that positive shocks to the state variable—which raise consumption volatility—are negatively correlated with consumption itself, so that high-volatility states coincide with low marginal-utility states for the investor. This sign restriction was motivated

[20] Both values are based on data in Shiller (2015). The former is the average growth of dividends over the period 1990–2024; the latter is the median of the standard deviations of dividend growth computed over non-overlapping five-year periods within the same sample. Both are reported across sample periods in Table 4 of Appendix B.

by He *et al.* (2017), who show that the intermediary capital ratio (the ratio of equity to assets, i.e. the inverse of leverage) is a priced state variable. When the capital ratio falls—equivalently, when leverage is high—the intermediary's portfolio value declines and risk premia rise. In the present consumption-based framework the capital ratio does not appear explicitly, but high leverage can be associated with high consumption volatility of the marginal investor, reflecting the amplified exposure of a leveraged intermediary to asset-price fluctuations. Requiring $\rho_{cx} < 0$ then adapts this mechanism. A deterioration in the intermediary's balance sheet (low capital ratio, high leverage) corresponds to a high-volatility state, and in that state consumption falls.

The initial parameter values supplied to the optimizer influence the solution found, and different starting points can yield different local optima. Nevertheless, high and strongly state-dependent consumption volatility arises consistently across the sweep: the spread $\sigma_c^{\mathrm{hi}} - \sigma_c^{\mathrm{lo}}$ is large in the baseline and robust across calibration runs, reflecting the structural need for volatility that is both high in level and variable across states in order to generate non-trivial risk premia.

Table 1: Baseline parameters calibrated by moment targeting. Parameters are chosen to minimize the sum of squared normalized deviations from the mean P/D ratio and the TIPS 10–5 yr yield spread. All other parameters are discussed in the text.

| Parameter | Symbol | Value |
|---|---|---|
| State variable persistence | $\phi$ | 0.8476 |
| Consumption–state correlation | $\rho_{cx}$ | $-0.6313$ |
| Consumption vol. (low state) | $\sigma_c^{\mathrm{lo}}$ | 0.2815 |
| Consumption vol. (high state) | $\sigma_c^{\mathrm{hi}}$ | 0.6323 |
| Risk-free rate (low state) | $r^{\mathrm{lo}}$ | 0.0000 |
| Risk-free rate (high state) | $r^{\mathrm{hi}}$ | 0.0996 |
| Time preference | $\rho$ | 0.0735 |

# 4 Results

This section presents the quantitative performance of the calibrated model and discusses the economic mechanisms behind the results. Table 2 reports the simulated moments against empirical targets. The empirical moments are drawn from two sources: stock-market moments (equity excess return, mean P/D ratio, median 2-year standard deviation of $\log(P/D)$, stock return volatility, and return predictability) are computed based on the monthly dataset of Shiller (2015) over

the period 1990–2024; the TIPS moments (10–5 yr yield spread and 10/5 yr excess return) are derived from the zero-coupon real yield curve estimated by Gürkaynak *et al.* (2010) over the period 2004–2024. The baseline directly targets two of these moments—the mean P/D ratio and the TIPS 10–5 yr yield spread—while the remaining moments are not targeted and are produced by the calibrated model. Figure 1 shows the key model-implied functions of the state variable $x$, together with the stationary distribution of $x$ and the decomposition of the 10-year real yield.

| Moment | Target | Model (SE) |
|---|---|---|
| Equity excess return ($\bar{r}_e$, %) | 8.384 | 5.387 (1.026) |
| Mean P/D ratio ($\overline{P/D}$) | 53.85 | 48.73 (4.766) |
| Median 2-yr std $\log(P/D)$ (med $\sigma_{pd}^{2y}$) | 0.0671 | 0.0562 (0.0112) |
| Std. stock return ($\sigma_e$, %) | 12.50 | 11.55 (2.244) |
| TIPS 10–5yr yield spread (TIPS sp., %) | 0.457 | 1.267 (0.112) |
| TIPS 10/5yr excess return (TIPS rx., %) | 0.708 | 2.188 (0.372) |
| Return pred. slope ($\hat{\beta}_1$) | 0.215 | 0.285 (0.128) |

Table 2: Simulated moments versus empirical targets. Stock-market moments are computed from the monthly dataset of Shiller (2015) over 1990–2024; TIPS moments are derived from the zero-coupon real yield curve of Gürkaynak *et al.* (2010) over 2004–2024. Returns and spreads are in %; the P/D ratio and $\hat{\beta}_1$ are dimensionless. Abbreviations in parentheses correspond to column headers in Table 7. The *Model* column reports the mean across 2,000 simulations; SE is the standard deviation of the moment estimate across separate simulations of the same length as the historical sample.

## 4.1 Consumption dynamics and the risk-free rate

Consumption volatility $\sigma_c(x)$ increases monotonically with the state variable, spanning more than a twofold range across the states that contain 95% of the stationary distribution. The unconditional average is approximately 0.44—substantially higher than aggregate consumption volatility for a representative household—consistent with the interpretation that the marginal investor is a leveraged intermediary whose consumption tracks asset prices. Consumption drift $\mu_c(x)$ also varies widely, increasing monotonically from low values at low $x$ to strongly positive values at high $x$.

The risk-free rate exhibits a minimum at a level of $x$ relatively close to the steady state, rising on both sides, though more steeply at high values of $x$ where consumption volatility is elevated. This U-shaped profile appears to be a feature

that this single-factor model requires in order to jointly reproduce the empirical moments considered here. A model with two state variables would offer more flexibility in shaping the interest-rate schedule. Importantly, this profile is not forced by the calibration: the values of the risk-free rate at the reference points $x = \pm 1$ are not explicitly targeted, and similar profiles arise across most of the 24 calibration runs reported in Appendix D. The model also does not generate a risk-free-rate puzzle, and its time-series variation is quite tame, as also illustrated by the sample simulation of the model in Appendix E.

## 4.2 Equity premium, excess volatility, and return predictability

The P/D ratio is targeted in the baseline, though the optimization algorithm is not able to match both targeted moments simultaneously with full precision, so the simulated P/D ratio lies somewhat below the empirical target of 53.9. Across all 24 calibration runs (Appendix D), the P/D ratio takes a wide range of values, and several runs produce levels comparable to those observed in the 1990–2024 sample as well as in earlier historical episodes when equity valuations were lower.[21]

The P/D ratio declines steeply with $x$: the stock market is highly valued when the state variable is low (consumption volatility low, discount rates low) and cheaply valued when $x$ is high. The expected equity excess return is correspondingly increasing in $x$: it is low when the state variable is low and rises as consumption volatility and the associated risk compensation increase. The equity excess return is not targeted, yet the model produces an unconditional mean of 5.4%—of the same order of magnitude as the 8.4% observed in the data. Across states, the conditional equity premium ranges from near zero in low-volatility environments to well above its unconditional average when the state variable is elevated, a range that seems broadly consistent with what is observed across different market regimes. The remaining gap relative to the empirical average could be closed by slightly higher risk aversion or by complementary mechanisms—such as rare disasters—without invalidating the consumption-volatility channel identified here.

The model naturally produces excess volatility in stock prices relative to dividends—the variance of stock prices exceeds what could be justified by the present value of future dividends computed with a constant discount rate, as

[21] In some runs the model becomes unstable and the P/D ratio diverges. This occurs when the effective discount rate under the dividend measure approaches zero, making the integral in Equation (17) non-convergent. These runs are identified in Appendix D.

documented by Shiller (1981). This excess volatility is a direct consequence of time-varying discount rates driven by the state variable: even with stable dividend fundamentals, stock prices fluctuate substantially. Quantitatively, the model generates a standard deviation of log real stock returns of 11.6%, close to the 12.5% observed in the 1990–2024 sample, and also closely matches the median 2-year rolling standard deviation of $\log(P/D)$ (0.056 in the model against an empirical value of 0.070; Table 4).[22]

Because the equity premium is state-dependent and the P/D ratio is a monotone function of the same state variable, current valuation ratios predict future returns, as documented empirically by Campbell and Shiller (1988). Table 3 reports the slopes from this predictive regression. The model implies slopes of 0.29, 0.79, and 1.01 at horizons of one, four, and eight years, consistent with the 1990–2024 estimates and with the well-documented pattern of predictability increasing with horizon.

Table 3: Return predictability regressions: OLS of the cumulative log real stock return $r_{t\to t+T}$ on the log dividend-price ratio $\log(D_t/P_t)$ at horizons $T = 1, 4, 8$ years, $r_{t\to t+T} = \alpha_T + \beta_T \log(D_t/P_t) + \varepsilon_{t+T}$. A positive slope $\beta_T$ indicates that a cheap market (low P/D ratio) predicts high future returns, consistent with time-varying discount rates. Empirical standard errors (in parentheses) are Newey–West with bandwidth equal to the horizon in months. Model standard errors are the cross-simulation standard deviation of $\hat{\beta}_T$. Model $R^2$ is the average in-sample $R^2$ across simulations.

| | Model | | | 1990–2024 | | | 1926–2024 | | |
|---|---|---|---|---|---|---|---|---|---|
| Horizon | $\beta$ | (s.e.) | $R^2$ | $\beta$ | (NW s.e.) | $R^2$ | $\beta$ | (NW s.e.) | $R^2$ |
| 1 year | 0.285 | (0.128) | 0.146 | 0.215 | (0.071) | 0.134 | 0.061 | (0.037) | 0.023 |
| 4 years | 0.791 | (0.338) | 0.400 | 0.831 | (0.230) | 0.454 | 0.209 | (0.096) | 0.079 |
| 8 years | 1.005 | (0.740) | 0.481 | 1.316 | (0.114) | 0.763 | 0.372 | (0.159) | 0.149 |

## 4.3 Bond term premia and the yield curve

Generating a positive and sizeable bond term premium is one of the central challenges of asset pricing models (Backus *et al.* 1989), and the challenge is even

[22]The full-sample standard deviation of $\log(P/D)$ in the data (0.27 over 1990–2024, up to 0.48 over 1926–2024) is likely inflated by slow-moving shifts in the long-run level of the P/D ratio—driven, for instance, by gradual changes in risk aversion, trend growth, or the equity risk premium over decades. Such level shifts would show up as a large standard deviation even if short-run fluctuations were modest. The 2-year rolling standard deviation captures local variability around whatever the prevailing level happens to be, and is therefore a cleaner target for a model with fixed structural parameters.

greater for real term premia. Long-run risk models, for example, can produce a high equity premium—using an inflation channel—but fail to deliver a significant real term premium, because the covariance of long-term bond returns with the stochastic discount factor is of the wrong sign or too small in magnitude.

In the present model, long-term bonds are risky. The risk-free rate rises sharply at high values of $x$, driven by consumption drift becoming strongly positive in high-volatility states. In bad states (high $x$), the marginal investor anticipates mean-reversion toward lower volatility and higher valuations. As rates rise, long-term bonds lose value, making them risky for the marginal investor and generating a positive real term premium. For this mechanism to operate, the increase in consumption drift must be large enough to overcome the precautionary savings effect, so that the risk-free rate rises on net. This sits in some tension with the standard intuition that distress periods are associated with low real interest rates. In a richer model with additional state variables, such an exact link between volatility and the interest rate would not be imposed. Understanding this relationship more precisely is an interesting avenue for further research, particularly given that real short-term rates are also shaped by monetary policy.

The yield decomposition shows that the 10-year real yield and the term premium both increase with $x$. At low values of $x$, both components are near zero or negative; in the central range of the stationary distribution the term premium is positive and increasing, reflecting that long-term bonds are risky precisely when the marginal investor's consumption volatility is high. At very high $x$ the total yield rises steeply, driven primarily by the risk-neutral yield (defined in Equation 12) as expected short rates become very high.

The bond moments used for comparison are the TIPS 10–5 yr yield spread and the TIPS 10/5 yr excess holding-period return (Table 5). The choice of spread and excess return rather than yield levels is deliberate. By focusing on the difference between two TIPS maturities, any liquidity premium that affects the overall level of TIPS yields cancels out. This should minimize the impact of liquidity effects in the data, which are not accounted for in the model, where all assets are assumed to be perfectly liquid. Similarly, restricting attention to TIPS avoids confounding real term premia with inflation risk premia that would be present in nominal bond yields.

The model generates a positive mean TIPS 10–5 yr yield spread and a positive mean TIPS 10/5 yr excess return. Notably, the model overshoots both moments relative to the empirical targets, even though one of them is directly targeted by the calibration. Despite the fact that generating a sizeable positive real term

premium is precisely what makes the bond premium puzzle difficult to solve, the consumption-volatility channel produces an excess even beyond what is observed in the data. Indeed, within this single-factor structure, the model requires a higher real term premium than the data in order to jointly match the remaining empirical moments. The yield decomposition into the risk-neutral yield ($RNY$, defined in Equation 12) and the term premium ($TP$, defined in Equation 11) closely mirrors the pattern documented for real yields in Abrahams *et al.* (2016), with the term premium accounting for a substantial and state-dependent share of the total yield.

Appendix F reports predictability regressions according to Fama and Bliss (1987) (FB) and Campbell and Shiller (1991) (CS). The model qualitatively violates the expectations hypothesis, even following the correct direction—FB slopes are positive and CS slopes are negative—but the estimates carry large standard deviations and do not fully reproduce the cross-sectional pattern of the empirical targets. These regressions should be interpreted with caution: the empirical targets were estimated for nominal bonds, and the model-implied slopes are not robust to moderate variations in the calibration. A more reliable indicator of the model's performance is that the functional forms of the yield curve and term premium are consistent with the decomposition of Abrahams *et al.* (2016).

Overall, the He *et al.* (2017) perspective—in which high consumption volatility coincides with distress and falling consumption, implying $\rho_{cx} < 0$—produces promising results across the range of moments considered here. Imposing the opposite sign restriction, consistent with the procyclical leverage documented by Adrian and Shin (2014), makes it considerably harder to match the empirical moments jointly, as the resulting calibration tends to generate risk premia of the wrong sign or insufficient magnitude. This does not mean that the Adrian and Shin (2014) perspective is inconsistent with the data. In a model with more than one state variable, the two channels could operate simultaneously, with different factors capturing different aspects of intermediary behavior, as explored in Kargar (2021). Alternatively, the procyclical-leverage mechanism may be more relevant for intermediaries that are primarily active in a specific asset class, rather than marginal investors across all asset classes simultaneously. In this case, matching moments across different asset classes would not be expected.

# 5 Recursive Preferences

The baseline model uses CRRA utility. This section shows that recursive utility, the framework of the well-known long-run risk models (Bansal and Yaron 2004),

does not generate positive real term premia, despite its success in explaining the equity premium.[23] Long-run risk models feature a time-varying consumption drift, and possibly a time-varying consumption diffusion. Since they use a representative agent, the drift and diffusion take values close to those observed in aggregate data.

Following Duffie and Epstein (1992), when utility is recursive, the utility specification is given by:

$$V_0 = \mathrm{E}_0 \int_0^\infty F(C_t, V_t)\mathrm{d}t \tag{19}$$

where $V_t$ is the value function at time $t$, and $F(\cdot,\cdot)$ determines flow utility. The latter is also the aggregator of the recursive utility process, and it is given by:

$$F(C_t, V_t) = \frac{\rho\,(1-\gamma)\,V_t}{1-1/\psi}\left(\left(\frac{C_t}{((1-\gamma)\,V_t)^{1/(1-\gamma)}}\right)^{1-1/\psi} - 1\right) \tag{20}$$

where $\psi$ is the intertemporal elasticity of substitution. The case of CRRA utility is nested in this specification, and it corresponds to the case where $\psi = 1/\gamma$. The SDF is derived from the expressions for the value function and the aggregator function. As shown by Tsai and Wachter (2018), the value function can be expressed as:[24]

$$V_t = \frac{C_t^{1-\gamma} e^{(1-\gamma)K(x_t)}}{1-\gamma} \tag{21}$$

where $V_t$ increases with $K$, which is a specific function of $x_t$ that captures the full dependence of the value function on the state variable.[25] Given the expression for the value function, Itô's lemma can be applied to derive the SDF process. Following Chen *et al.* (2009), the fundamental relationship is:

$$\frac{\mathrm{d}\Lambda}{\Lambda} = F_V(C_t, V_t)\mathrm{d}t + \frac{\mathrm{d}F_C(C_t, V_t)}{F_C(C_t, V_t)} \tag{22}$$

$F_C$ and $F_V$ denote partial derivatives of $F$ with respect to consumption and the value function respectively. The first term on the right-hand side is the derivative of the flow utility with respect to the value function. The second term is obtained by

[23] As has been mentioned before, an earlier version of this article goes through many variations that show it is not possible to generate high and positive bond term premia.

[24] Similar results are common in the literature, see for example Benzoni *et al.* 2011; Kraft *et al.* 2017.

[25] The model is not solved fully in this section because this is not required for the analysis. However, the earlier version of the article contains a full solution of the model. In addition, Melissinos (2023a) provides a solution method based on a perturbation expansion of $K(\cdot)$.

applying Itô's lemma to the derivative of flow utility with respect to consumption:[26]

$$\begin{aligned}\frac{\mathrm{d}\Lambda}{\Lambda} =&\left(\frac{\rho(-(1-\gamma\psi)e^{\frac{(1-\psi)K[x_t]}{\psi}}-\gamma\psi+\psi)}{1-\psi}-\gamma\mu_{ct}+\frac{\gamma^2\sigma_{ct}^2}{2}+\frac{\gamma(\gamma\psi-1)\rho_{cxt}\sigma_{xt}\sigma_{ct}K'(x_t)}{\psi}\right.\\&\left.+\frac{(\gamma\psi-1)(-2\psi\log(\phi)x_tK'(x_t)+\sigma_{xt}^2((\gamma\psi-1)K'(x_t)^2-\psi K''(x_t)))}{2\psi^2}\right)\mathrm{d}t\\&-\frac{(\gamma\psi-1)\sigma_{xt}K'(x_t)}{\psi}\mathrm{d}W_{xt}-\gamma\sigma_{ct}\mathrm{d}W_{ct}\end{aligned}\tag{23}$$

When $\gamma = 1/\psi$ (time-separable utility), the equation reduces to Equation 5. Moreover, the consumption-driven stochastic component ($-\gamma\sigma_{ct}\mathrm{d}W_{ct}$) is identical to the time-separable case; the extra term $-\frac{(\gamma\psi-1)\sigma_{xt}K'(x_t)}{\psi}\mathrm{d}W_{xt}$ arises from the explicit dependence of the SDF on the state variable.

Using the SDF expression, the risk-free rate, stock price, and bond price can be derived as before. The analysis focuses on the bond to show that recursive utility generates a negative real term premium. Specifically, the expected excess return of the bond can be expressed as:

$$\underbrace{-\frac{B_m}{B}+\log(\phi)x_t\frac{B_x}{B}+\frac{\sigma_{xt}^2}{2}\frac{B_{xx}}{B}-r(x_t)}_{\text{derived from: } \mathrm{E}[\frac{\mathrm{d}B}{B}]-r\mathrm{d}t}=\left(\rho_{cxt}\gamma\sigma_{ct}\sigma_{xt}+\frac{(\gamma\psi-1)\sigma_{xt}^2K'(x_t)}{\psi}\right)\frac{B_x}{B}\tag{24}$$

The second term in the parenthesis is unique to recursive utility. Furthermore, this term is in general negative. Whether $x_t$ represents the consumption drift or its diffusion, the following qualitative relationships typically hold in equilibrium:

$$\left.\begin{aligned}&(x\uparrow\Rightarrow\mu_c\uparrow)\Rightarrow(K'(x)>0)\text{ and }(r'(x)>0\Rightarrow B_x<0)\\&(x\uparrow\Rightarrow\sigma_c\uparrow)\Rightarrow(K'(x)<0)\text{ and }(r'(x)<0\Rightarrow B_x>0)\end{aligned}\right\}\Rightarrow K'(x)B_x<0\tag{25}$$

The first case corresponds to a growth shock: an increase in expected growth improves welfare ($K'(x) > 0$) but raises interest rates, causing bond prices to fall ($B_x < 0$). The second case corresponds to a volatility shock: an increase in uncertainty reduces welfare ($K'(x) < 0$) but triggers precautionary savings, lowering rates and raising bond prices ($B_x > 0$). In both scenarios, the product $K'(x)B_x$ is negative. Given a preference for early resolution of uncertainty ($\gamma > 1/\psi \Rightarrow \gamma\psi - 1 > 0$), the recursive utility component in Equation (24) acts as a "hedging"

[26]This operation is performed by substituting the value function using Equation (21) and applying Itô's lemma based on consumption and the state variable as independent variables.

term. This makes bonds an attractive asset for an investor worried about future states, thereby pushing the real risk premium downward.[27] Consequently, while recursive utility can be a powerful tool for generating a high equity premium, it tends to produce negative real term premia. This is also consistent with Bansal and Shaliastovich (2013), where recursive utility is used in a long-run risk model, and a positive nominal term premium is generated via an inflation channel. However, the underlying real term premium is still negative.

# 6 Discussion

A natural question is whether the high consumption volatility ascribed to the marginal investor is empirically plausible. This is not straightforward to answer, in part because it is difficult to identify who the marginal investor actually is. Nevertheless, even outside the intermediary literature, wealthy individuals who are more likely to be the marginal investors have substantially more volatile consumption than median households (Ait-Sahalia *et al.* 2004).

A second, perhaps even more natural, identification is with financial intermediaries, which is the primary perspective of this article. The performance of intermediaries is directly tied to the value of leveraged portfolios of financial assets. Because financial asset prices are themselves highly volatile, and because leverage amplifies this volatility, it is not surprising that intermediary consumption volatility would substantially exceed that of the aggregate household sector. This is precisely the mechanism formalized by He *et al.* (2017).

Despite the difficulty of directly observing intermediary consumption, Appendix G provides indirect evidence for the high volatility and its connection to asset prices. Using daily equity return data for NY Fed primary dealer holding companies from He *et al.* (2017), a 1-year rolling volatility series is interpreted as the consumption diffusion $\sigma_c(x)$ and then inverted to recover the state variable $x_t$ at each date. This produces a time series of the implied state variable and the corresponding model-implied asset prices. Two findings stand out. First, while the sample mean of intermediary equity return volatility (23%) lies below the model's

[27]This result is reinforced by solving the model using several calibrations, in an earlier version of this article Melissinos (2023b). In all cases the contribution to the term premium was negative. A fully general proof is not provided. The complication can arise from the fact that the risk-free rate is not only a function of consumption drift and diffusion, but also a function of the recursive utility term. In addition, the state variable can affect both consumption drift and consumption diffusion at the same time and these relationship in theory could be arbitrarily non-linear. In these cases, it becomes more difficult to unambiguously determine the sign of $K'(x)B_x$.

unconditional mean $\sigma_c$ of 44%, both figures are orders of magnitude above standard aggregate consumption volatility, confirming that the high-volatility premise of the model is structurally supported by the intermediary data. In addition, during stress episodes—most notably the 2008 financial crisis—the series reaches values above 70%, exceeding even the model's high-volatility states. Second, the model-implied P/D ratio exhibits a positive correlation of 0.48 in levels and 0.35 in changes with its empirical counterpart, and the model-implied 10–5 yr TIPS spread displays a positive levels correlation of 0.33. These correlations are moderate rather than large, unsurprisingly, given that the model has only one state variable. Nevertheless, the evidence suggests that the state variable $x_t$ captures a meaningful fraction of the common variation in asset prices, consistent with the interpretation that intermediary distress is a key driver of risk premia.

It is worth stressing why high consumption volatility is the key ingredient. Based on Section 5 and the model variations explored in Melissinos (2023b), matching the main empirical moments relating to risk premia of stocks and bonds simultaneously does not appear possible with standard preference specifications and a single state variable. Even recursive utility (long-run risk) models can match the equity premium, but by themselves they cannot generate positive real bond premia. High and time-varying consumption volatility of the marginal investor—far above what aggregate data suggest—is therefore not merely one possible mechanism but an essential ingredient.

This model is, of course, highly stylized: it features a single state variable and abstracts from many features of actual financial markets. Yet precisely because of this parsimony, it is striking how naturally it accommodates several classic asset pricing puzzles simultaneously—the equity premium, excess return volatility, return predictability, and positive real term premia—within a unified framework. Richer models with additional state variables could plausibly improve the quantitative fit, but the present results suggest that high and time-varying consumption volatility of the marginal investor is a parsimonious source of the various asset pricing patterns documented here.

# 7 Conclusion

This article has shown that a parsimonious single-factor model, in which the marginal investor faces high and time-varying consumption volatility consistent with the intermediary asset pricing perspective of He *et al.* (2017), is able to jointly address several classic asset pricing puzzles, including the equity premium puzzle,

the risk-free rate puzzle, excess return volatility, return predictability, and—most challenging—a positive and time-varying real term premium. Addressing these puzzles does not mean that every empirical moment is precisely matched; rather, the model reproduces the qualitative behavior that standard frameworks cannot. For instance, it generates a positive and sizable real term premium where those frameworks predict a near-zero or negative one.

In retrospect, the resolution may have been staring us in the face all along. The equity premium puzzle has long signaled, within the representative-agent framework, that consumption volatility would need to be far higher than what aggregate data suggest to rationalize observed risk premia. Rather than a symptom of model failure, this can be taken as a hint that the consumption volatility of the marginal investor is genuinely high and time-varying. For regular households, whose consumption is relatively stable, the high risk premia of financial markets represent returns that far exceed what their own risk profile would demand. This suggests that limited participation in financial markets by regular households (or other institutions not exposed to high volatility) may represent a significant missed opportunity, with implications for wealth accumulation and the distribution of asset returns across the population. In addition, if median households were to participate more actively in equity and bond markets, the resulting shift in the marginal investor's risk profile would likely compress both the equity premium and the real term premium.

The results suggest several directions for further research. A natural first extension is to introduce a second state variable. The single-factor structure generates an inherent tension between the equity premium and the real term premium: the same consumption-volatility channel that raises the equity premium pushes the term premium higher than observed in the data. A second factor could separately discipline the two risk premia, potentially resolving this tension and improving the quantitative fit of the bond moments without sacrificing the equity results. A second state variable could also capture time variation in the risk-free rate that is orthogonal to the volatility cycle, such as a slow-moving component of consumption drift associated with long-run growth expectations. Part of the challenge in this direction is to identify all the parameters that govern the joint process of the two state variables.

A second direction concerns the use of the model as a lens for understanding historical asset price episodes. The indirect evidence in Appendix G already suggests that the implied state variable $x_t$ comoves meaningfully with empirical P/D ratios and TIPS spreads. A more systematic historical decomposition—using

the inversion approach to back out $x_t$ and then attributing observed movements in yields and equity valuations to the model's consumption-volatility mechanism—could shed light on episodes such as the 2008 financial crisis, the post-crisis period of compressed term premia, and the 2022 rate cycle.

Finally, in a more extensive framework, one could move beyond the exogenous consumption process and embed the marginal investor in a general equilibrium framework. Doing so would allow one to endogenize the consumption dynamics of the intermediary sector, connecting the reduced-form process to observable balance-sheet quantities such as leverage and capital ratios, as in He and Krishnamurthy (2013). It would also make it possible to study welfare: how do movements in risk premia redistribute between the intermediary sector and ordinary households, and what are the implications for the distributional effects of monetary policy?

# A Feynman-Kac Formula

The differential equations arising from the pricing equations for the bond and the strip price-dividend ratio are solved using the Feynman-Kac formula.

## A.1 Feynman-Kac Formula for the Bond Price

Equation 9 is repeated here:

$$-B_m - r(x_t)B + \underbrace{\left(\log(\phi)x_t - \rho_{cxt}\gamma\sigma_{ct}\sigma_{xt}\right)}_{\tilde{\mu}(x_t)} B_x + \frac{\sigma_{xt}^2}{2}B_{xx} = 0$$

The Feynman-Kac formula states that the solution to this equation is given by:

$$B(x_t, m) = \mathrm{E}\left[\exp\left(-\int_0^m r(\tilde{x}_{t+s})\mathrm{d}s\right)\Big|\tilde{x}_t = x_t\right] \tag{26}$$

$$\mathrm{d}\tilde{x}_t = \tilde{\mu}(\tilde{x}_t)\mathrm{d}t + \sigma_{xt}(\tilde{x}_t)\mathrm{d}W_{xt} \tag{27}$$

where $\tilde{\mu}(\cdot)$ is defined by the term multiplying $B_x$ above.

## A.2 Feynman-Kac Formula for the Strip Price-Dividend Ratio

Equation 15 is repeated here:

$$-\hat{s}_m - \underbrace{\left(r(x) - \mu_D + \rho_{cD}\gamma\sigma_{ct}\sigma_D\right)}_{\hat{r}(x_t)}\hat{s} + \underbrace{\left(\log(\phi)x_t - \rho_{cxt}\gamma\sigma_{ct}\sigma_{xt} + \rho_{xD}\sigma_{xt}\sigma_D\right)}_{\hat{\mu}(x_t)}\hat{s}_x + \frac{\sigma_{xt}^2}{2}\hat{s}_{xx} = 0$$

The Feynman-Kac formula states that the solution to this equation is given by:

$$\hat{s}(x_t, m) = \mathrm{E}\left[\exp\left(-\int_0^m \hat{r}(\hat{x}_{t+s})\mathrm{d}s\right)\Big|\hat{x}_t = x_t\right] \tag{28}$$

$$\mathrm{d}\hat{x}_t = \hat{\mu}(\hat{x}_t)\mathrm{d}t + \sigma_{xt}(\hat{x}_t)\mathrm{d}W_{xt} \tag{29}$$

where $\hat{r}(\cdot)$ is defined by the term multiplying $\hat{s}$ above and $\hat{\mu}(\cdot)$ is defined by the term multiplying $\hat{s}_x$ above.

# B Empirical Moments

The main empirical moments are reported in Tables 2 and 3 in the main text. Tables 4 and 5 below extend these by reporting moments across multiple sample periods for comparison. For the TIPS data, the moments from the 2005–2024 subsample, when TIPS were more liquid, are close to those from the full 1999–2024 sample, suggesting that liquidity conditions do not materially affect the empirical targets.

Table 4: Stock market and macroeconomic moments: model vs. data. Empirical data: Shiller (prices, dividends, CPI, total return index) merged with Fama–French monthly risk-free rate. All returns and rates are annualized (% unless noted).

| Moment | Model | 1990–2024 | 1999–2024 | 1926–2024 |
|---|---|---|---|---|
| Mean P/D ratio | 48.7 | 53.85 | 57.97 | 34.51 |
| Median 2-yr std $\log(P/D)$ | 0.056 | 0.07 | 0.07 | 0.10 |
| Mean real div. growth (%) | — | 2.89 | 3.41 | 1.97 |
| Std. real div. growth (%) | — | 2.47 | 2.78 | 4.02 |
| Mean real stock return (%) | 4.56 | 8.42 | 6.20 | 8.16 |
| Std. log real stock return (%) | 11.55 | 12.50 | 13.08 | 15.30 |
| Mean equity premium (%) | 5.39 | 8.38 | 6.84 | 7.86 |

Table 5: TIPS term structure moments: model vs. data. Source: Federal Reserve FEDS 2008-05 zero-coupon TIPS yield curve. Yields are continuously compounded. Excess returns are monthly log holding-period returns, annualized ($\times 12$). Fractional-maturity yields via Nelson–Siegel–Svensson (GSW 2008).

| Moment (%) | Model | 1999–2024 | 2005–2024 |
|---|---|---|---|
| Mean 10yr$-$5yr yield spread | 1.267 | 0.457 | 0.440 |
| Mean excess return 10yr/5yr | 2.188 | 0.708 | 0.576 |

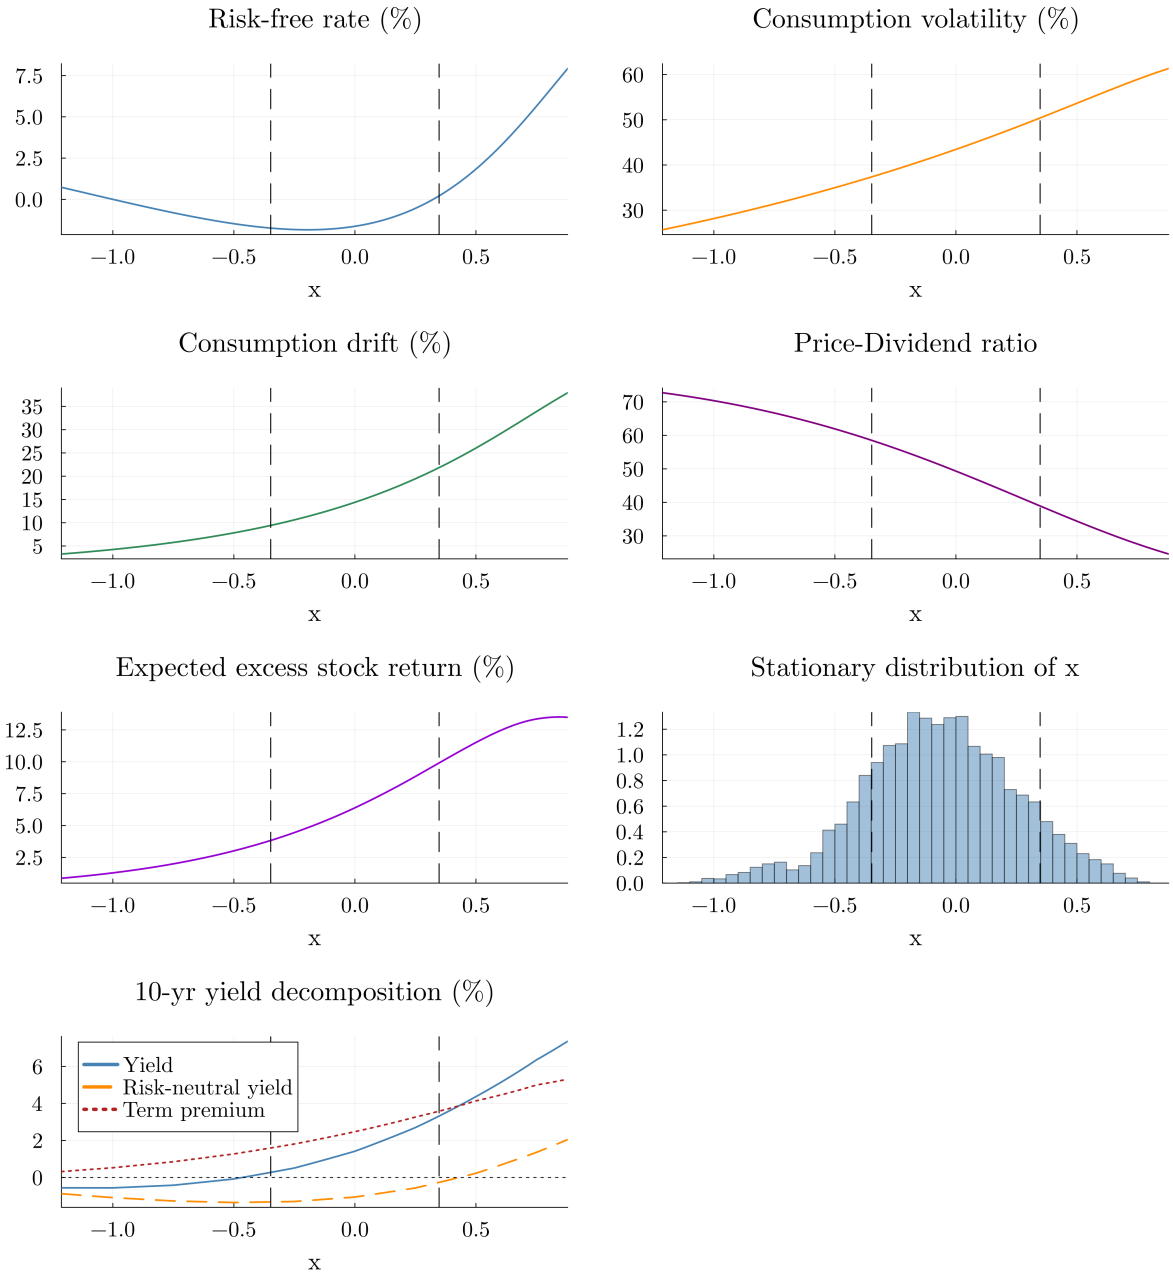


Figure 1: Model-implied quantities under the baseline calibration. Risk-free rate $r(x)$; consumption volatility $\sigma_c(x)$; consumption drift $\mu_c(x)$; price-dividend ratio $P/D(x)$; expected equity excess return $\mathbb{E}[R_e - r](x)$; stationary distribution of $x$; decomposition of the 10-year real bond yield into the risk-neutral yield and the term premium. Dashed vertical lines mark one stationary standard deviation from the steady state.

# C Functional Forms for Consumption Dynamics

As mentioned in the main text, log consumption follows the continuous-time diffusion

$$d\log C_t = \mu_c(x_t)\,dt + \sigma_c(x_t)\,dW_t^c, \tag{30}$$

where $x_t$ is a scalar latent state variable and $W^c$ is a standard Brownian motion. Both the drift $\mu_c$ and the diffusion $\sigma_c$ belong to the exponential family

$$\sigma_c(x) = \sigma_{c0}\,\exp\big(\sigma_{c1}\,s(x)\big), \tag{31}$$

$$\mu_c(x) = \mu_{c0}\,\exp\big(\mu_{c1}\,s(x)\big), \tag{32}$$

with four shape parameters ($\sigma_{c0}$, $\sigma_{c1}$, $\mu_{c0}$, $\mu_{c1}$). When $\sigma_{c1} > 0$ and $\mu_{c1} > 0$, both functions are strictly increasing in $x$: the high-$x$ regime is a high-volatility, high-growth state and the low-$x$ regime is a low-volatility, low-growth state.

The argument $s(x)$ is a smooth soft-cap transformation,

$$s(x) = 1 - \frac{1}{k}\log\Big(1 + e^{k(1-x)}\Big), \qquad k = 5, \tag{33}$$

which satisfies $s(x) \approx x$ for $x \ll 1$ (so the functions behave like standard exponentials in the range the state variable typically occupies) and $s(x) \to 1$ as $x \to +\infty$ (so $\sigma_c$ and $\mu_c$ are asymptotically bounded above by $\sigma_{c0}\,e^{\sigma_{c1}}$ and $\mu_{c0}\,e^{\mu_{c1}}$ respectively; note that $s(1) = 1 - \frac{\log 2}{k} < 1$, so the asymptotic upper bound is strictly greater than the value at $x = 1$). This cap prevents the dividend-measure drift from becoming explosive at extreme positive values of the state variable, which would cause the P/D integral to diverge. There is no analogous cap on the negative side: as $x \to -\infty$, $s(x) \to x$, so $\sigma_c$ and $\mu_c$ decay exponentially, and the economy approaches a deterministic low-growth steady state.

The parameters that determine the functional forms of $\sigma_c$ and $\mu_c$ are $\sigma_{c0}$, $\sigma_{c1}$, $\mu_{c0}$, and $\mu_{c1}$. These are chosen according to the process described in the Subsection D of the Appendix.

# D SMM Sweep

As described in the main text, some parameters are directly calibrated. This includes the drift and diffusion of the dividend process, the risk aversion parameter, $\gamma$, and the diffusion parameter of the state variable, $\sigma_x$. In addition, the steady state level of the state variable, $\bar{x}$, is set to zero without loss of generality, and the

correlation parameter between consumption and dividend shocks, $\rho_{cD}$, is also set to 0. The remaining free parameters are chosen to match empirical moments.

The main text uses a single baseline parameterization (run 11). The full sweep of 24 runs is reported here to illustrate the robustness of the qualitative results. The individual moment values are sensitive to the choice of targeted moments and to the initialization of the optimizer—as discussed further below, the starting point $\theta_0$ does influence the solution found by the Nelder–Mead algorithm. Nevertheless, across all 24 runs a non-negligible equity premium, excess stock-return volatility, return predictability, and a positive long-term bond risk premium invariably arise.[28]. The one structural assumption that is imposed throughout—and that is the primary driver of the sign and direction of all risk premia—is the sign of $\rho_{cx}$, which is restricted to be negative. As discussed in the main text, this sign determines whether the price of risk with respect to consumption volatility is positive or negative, and hence whether equity and bond premia are positive.

## D.1 Calibration approach

The model contains seven free parameters: the persistence of the state variable $\phi$; its correlation with consumption shocks $\rho_{cx}$; the consumption-volatility schedule evaluated at the two boundary states, $\sigma_c^{\text{lo}} \equiv \sigma_c(x{=}-1)$ and $\sigma_c^{\text{hi}} \equiv \sigma_c(x{=}+1)$; the analogous risk-free rate schedule $r^{\text{lo}}$ and $r^{\text{hi}}$; and the time-preference rate $\rho$. The remaining structural parameter, $\rho_{cD}$ (the correlation between consumption and dividend shocks), is fixed at zero throughout.[29]

Seven empirical moments are used as candidate calibration targets:

- The mean equity excess return $\bar{r}_e$ (Shiller data, 1990–2024, annualized);
- The median 2-year rolling standard deviation of $\log(P/D)$ (Shiller data, 1990–2024);
- The standard deviation of log real stock returns $\sigma_e$ (Shiller data, 1990–2024);
- The mean P/D ratio $\overline{P/D}$ (Shiller data, 1990–2024);
- The mean TIPS 10–5 yr yield spread (Federal Reserve FEDS 2008-05 zero-coupon TIPS curve, 2004–2024);

[28]Some runs yield unstable or extreme simulated moments—for instance, a very large mean P/D ratio—indicating that certain parameter combinations push the model outside a well-behaved region of the parameter space

[29]Fixing $\rho_{cD} = 0$ imposes no constraint on the equity risk premium because the dividend shock affects asset prices only through the P/D interpolant, not through the SDF.

- The mean TIPS 10/5 yr log excess holding-period return (same source and sample);
- The slope coefficient from a univariate regression of cumulative log stock returns on $\log(D/P)$ at a one-year horizon, $\hat{\beta}_1$ (Shiller data, 1990–2024).

For a given subset $\mathcal{M}$ of these moments, the calibration objective is

$$\mathcal{L}(\theta) = \sum_{m \in \mathcal{M}} \left( \frac{\hat{m}(\theta) - m^*}{m^*} \right)^2, \tag{34}$$

where $\hat{m}(\theta)$ is the simulated moment at parameters $\theta$ and $m^*$ is the corresponding empirical target. Moments are weighted equally after normalization by the target value; no formal optimal- weighting procedure is used.

Optimization uses the Nelder–Mead simplex algorithm (Mogensen and Riseth, 2018) applied to logit-transformed parameters, which maps each bounded parameter onto the real line and prevents the optimizer from leaving the admissible region. Bounds are $\phi \in [0.40, 0.96]$, $\rho_{cx} \in [-0.90, -0.10]$, $\sigma_c^{\text{lo}} \in [0.05, 0.45]$, $\sigma_c^{\text{hi}} \in [0.06, 0.65]$, $r^{\text{lo}} \in [0.00, 0.02]$, $r^{\text{hi}} \in [0.03, 0.10]$, and $\rho \in [0.001, 0.100]$. Each run proceeds in two stages: a coarse stage with 20 simulation paths per moment evaluation and 500 Nelder–Mead iterations, followed by a medium-resolution stage with 40 paths and 300 iterations starting from the best coarse solution. All 24 runs are initialized from the same starting point $\theta_0$, shown in the caption of Table 6.

## D.2 Calibration of shape parameters

The four shape parameters ($\sigma_{c0}$, $\sigma_{c1}$, $\mu_{c0}$, $\mu_{c1}$) are not estimated directly. Instead, the SMM treats as free parameters the boundary evaluations

$$\sigma_c^{\text{lo}} \equiv \sigma_c(-1), \quad \sigma_c^{\text{hi}} \equiv \sigma_c(+1), \quad r^{\text{lo}} \equiv r(-1), \quad r^{\text{hi}} \equiv r(+1), \tag{35}$$

and recovers the shape parameters by solving the implied $4 \times 4$ nonlinear system. Concretely, given $\sigma_c^{\text{lo}}$ and $\sigma_c^{\text{hi}}$, equations (31) evaluated at $x = \pm 1$ form a $2 \times 2$ system in $(\sigma_{c0}, \sigma_{c1})$, which is solved first. Given the solution, equations (6) and (32) evaluated at $x = \pm 1$ form a $2 \times 2$ system in $(\mu_{c0}, \mu_{c1})$, which is solved second. Both systems admit closed-form analytical solutions: dividing the two equations in each system and taking logarithms isolates the shape exponent, after which the scale parameter follows directly.[30]

[30]The sequential structure avoids a four-dimensional root-find: the volatility shape parameters do not appear in the mean-growth system, so the two $2 \times 2$ systems can be solved in sequence.

This reparametrization has two advantages. First, the boundary values ($\sigma_c^{\text{lo}}$, $\sigma_c^{\text{hi}}$, $r^{\text{lo}}$, $r^{\text{hi}}$) have direct economic interpretations—they pin down the level and slope of the volatility and interest-rate schedules—and can be bounded naturally in the SMM. Second, the shape parameters ($\sigma_{c0}$, $\sigma_{c1}$, $\mu_{c0}$, $\mu_{c1}$) are fully determined by the boundary values, so the number of free parameters does not increase with the richness of the functional form.

Table 6 in Appendix D reports the calibrated boundary values ($\sigma_c^{\text{lo}}$, $\sigma_c^{\text{hi}}$, $r^{\text{lo}}$, $r^{\text{hi}}$) for all 24 sweep runs.

## D.3 The sweep design

Because the seven candidate moments are not all jointly attainable at a single parameter vector—the model faces structural tensions between matching a high mean P/D ratio and a large equity premium, and between matching TIPS term-structure moments and equity moments simultaneously— a systematic sweep is run that varies which moments enter the objective function. This traces out the Pareto frontier of moment fit and identifies which combinations can be achieved simultaneously.

The 24 runs are organized into four groups:

1. **Group 1 (individual moments, 7 runs).** Each run targets a single moment in isolation. When a moment can be matched individually, the objective typically converges to near zero, confirming that the model is flexible enough to match that moment on its own.

2. **Group 2 (pairs anchored on the TIPS spread, 6 runs).** Each run targets the TIPS 10–5 yr yield spread together with one other moment. The TIPS spread is held as an anchor because it is a clean real-rate term-structure moment that disciplines the level of the yield curve independently of the equity premium.

3. **Group 3 (pairs anchored on the mean P/D ratio, 6 runs).** Each run targets the mean P/D ratio together with one other moment. These runs quantify the trade-off between anchoring valuation ratios at their historical level and matching return moments.

4. **Group 4 (triples with both anchors, 5 runs).** Each run targets both the TIPS spread and the mean P/D ratio simultaneously, together with one additional moment. These runs represent the most demanding joint requirements.

## D.4 High-quality moment re-evaluation

The objective (34) is evaluated with a coarse simulation ($n = 20$–$40$ paths) during optimization to limit runtime. To obtain reliable estimates of what each parameterization actually generates, all 24 calibrated parameter vectors are subsequently re-evaluated using $n = 2{,}000$ simulation paths of length $T = 35$ years each, with 200 Monte Carlo trajectories for the bond and stock price solvers. These high-quality estimates are reported in Table 7.

A notable implication of this re-evaluation concerns the mean P/D ratio. Because the stationary P/D distribution is right-skewed—the ratio is bounded below at zero but can be very large in states where the effective discount rate under the dividend measure is low—the sample mean based on only $n = 20$–$40$ paths has a systematic upward bias. As a result, the mean P/D values shown in Table 7 are somewhat lower than the in-sample estimates seen by the optimizer: the high-quality re-evaluation reveals that the Group 3 and Group 4 parameterizations generate a mean P/D of approximately 46–50, compared to the nominal target of 53.85 that the optimizer nominally matched. The $n = 2{,}000$ estimates are the ones used throughout the article.

## D.5 Results

Table 6 reports the calibrated parameter vector and SMM objective value for all 24 runs. Table 7 reports the corresponding simulated moments, with the moments targeted in each run shown in bold.

Table 6: SMM sweep: calibrated parameters for all 24 runs. All 24 runs shared the same starting point $\theta_0$ (shown in the first row).

| Run | Targeted moments | $\phi$ | $\rho_{cx}$ | $\sigma_c^{\text{lo}}$ | $\sigma_c^{\text{hi}}$ | $r^{\text{lo}}$ | $r^{\text{hi}}$ | $\rho$ | Obj. |
|---|---|---|---|---|---|---|---|---|---|
| *Starting point $\theta_0$* | | | | | | | | | |
| $\theta_0$ | – | 0.847 | -0.630 | 0.280 | 0.617 | 0.000 | 0.100 | 0.074 | – |
| *Group 1: individual moment (one target each)* | | | | | | | | | |
| 1 | eq. ret. | 0.9047 | -0.6917 | 0.2713 | 0.6415 | 0.0000 | 0.0994 | 0.0655 | 0.00e+00 |
| 2 | med. 2yr $\sigma_{pd}$ | 0.8976 | -0.6337 | 0.2722 | 0.6408 | 0.0000 | 0.0999 | 0.0757 | 5.00e-06 |
| 3 | std. ret. | 0.8716 | -0.6446 | 0.2779 | 0.6083 | 0.0000 | 0.0999 | 0.0775 | 2.60e-05 |
| 4 | P/D ratio | 0.8490 | -0.6358 | 0.2834 | 0.6381 | 0.0000 | 0.0992 | 0.0749 | 4.0125 |
| 5 | TIPS sp. | 0.6989 | -0.6228 | 0.2968 | 0.6284 | 0.0000 | 0.0999 | 0.0816 | 0.00e+00 |
| 6 | TIPS rx. | 0.7289 | -0.6040 | 0.3070 | 0.6094 | 0.0000 | 0.0999 | 0.0765 | 9.00e-06 |
| 7 | pred. $\hat{\beta}_1$ | 0.8638 | -0.4638 | 0.2996 | 0.3927 | 0.0000 | 0.0993 | 0.0853 | 9.20e-05 |
| *Group 2: pairs anchored on TIPS 10–5yr spread* | | | | | | | | | |
| 8 | eq. ret., TIPS sp. | 0.8341 | -0.6097 | 0.2263 | 0.5080 | 0.0000 | 0.0926 | 0.0902 | 0.00e+00 |
| 9 | med. 2yr $\sigma_{pd}$, TIPS sp. | 0.9156 | -0.3388 | 0.2918 | 0.6276 | 0.0000 | 0.1000 | 0.0701 | 6.10e-05 |
| 10 | std. ret., TIPS sp. | 0.9381 | -0.5457 | 0.2369 | 0.3085 | 0.0000 | 0.1000 | 0.0847 | 5.30e-05 |
| 11 | P/D ratio, TIPS sp. | 0.8476 | -0.6313 | 0.2815 | 0.6323 | 0.0000 | 0.0996 | 0.0735 | 21.8529 |
| 12 | TIPS sp., TIPS rx. | 0.7467 | -0.5587 | 0.2691 | 0.6340 | 0.0000 | 0.0734 | 0.0851 | 8.00e-06 |
| 13 | TIPS sp., pred. $\hat{\beta}_1$ | 0.9554 | -0.4741 | 0.2625 | 0.3600 | 0.0000 | 0.0902 | 0.0851 | 0.0035 |

*continued on next page*

Table 6 (continued)

| Run | Targeted moments | $\phi$ | $\rho_{cx}$ | $\sigma_c^{\text{lo}}$ | $\sigma_c^{\text{hi}}$ | $r^{\text{lo}}$ | $r^{\text{hi}}$ | $\rho$ | Obj. |
|---|---|---|---|---|---|---|---|---|---|
| *Group 3: pairs anchored on mean P/D ratio* | | | | | | | | | |
| 14 | eq. ret., P/D ratio | 0.8490 | -0.6362 | 0.2828 | 0.6472 | 0.0000 | 0.0988 | 0.0746 | 4.3513 |
| 15 | P/D ratio, med. 2yr $\sigma_{pd}$ | 0.8491 | -0.6346 | 0.2830 | 0.6457 | 0.0000 | 0.0992 | 0.0747 | 4.2819 |
| 16 | P/D ratio, std. ret. | 0.8487 | -0.6378 | 0.2825 | 0.6457 | 0.0000 | 0.0991 | 0.0748 | 4.3507 |
| 17 | P/D ratio, TIPS sp. | 0.8476 | -0.6313 | 0.2815 | 0.6323 | 0.0000 | 0.0996 | 0.0735 | 21.8529 |
| 18 | P/D ratio, TIPS rx. | 0.8480 | -0.6257 | 0.2850 | 0.6058 | 0.0000 | 0.0997 | 0.0722 | 2.5849 |
| 19 | P/D ratio, pred. $\hat{\beta}_1$ | 0.8509 | -0.6298 | 0.2835 | 0.6191 | 0.0000 | 0.0998 | 0.0752 | 3.4350 |
| *Group 4: triples anchored on TIPS spread + P/D ratio* | | | | | | | | | |
| 20 | eq. ret., P/D ratio, TIPS sp. | 0.8490 | -0.6359 | 0.2828 | 0.6469 | 0.0000 | 0.0989 | 0.0746 | 4.3458 |
| 21 | P/D ratio, med. 2yr $\sigma_{pd}$, TIPS sp. | 0.8476 | -0.6313 | 0.2815 | 0.6323 | 0.0000 | 0.0996 | 0.0735 | 21.8533 |
| 22 | P/D ratio, std. ret., TIPS sp. | 0.8486 | -0.6362 | 0.2828 | 0.6454 | 0.0000 | 0.0990 | 0.0743 | 4.2243 |
| 23 | P/D ratio, TIPS sp., TIPS rx. | 0.8485 | -0.6324 | 0.2855 | 0.6084 | 0.0000 | 0.0990 | 0.0745 | 2.8480 |
| 24 | P/D ratio, TIPS sp., pred. $\hat{\beta}_1$ | 0.8515 | -0.6254 | 0.2852 | 0.6197 | 0.0000 | 0.0997 | 0.0758 | 3.3985 |

Table 7: SMM sweep: simulated moments for all 24 runs ($n$ = 2,000 simulations, $T$ = 35 years per path). Returns, spreads, and rates are in %; P/D ratio and $\hat{\beta}_1$ are in levels. **Bold** values indicate that the moment was included as a target in the SMM objective for that run. *n.c.* = not converged: the dividend-strip integral did not decay to less than 1% of the total at the 200-year horizon, so stock-market moments are unreliable for that run. Bond moments (TIPS spread, TIPS excess return) are unaffected.

| Run | Targeted moments | $\bar{r}_e$ | $\overline{P/D}$ | med $\sigma_{pd}^{2y}$ | $\sigma_e$ | TIPS sp. | TIPS rx. | $\hat{\beta}_1$ |
|---|---|---|---|---|---|---|---|---|
| Empirical target | | 8.384 | 53.8 | 0.0671 | 12.497 | 0.457 | 0.708 | 0.2152 |
| *Group 1: individual moment (one target each)* | | | | | | | | |
| 1 | eq. ret. | **7.189** | 23.9 | 0.0693 | 13.864 | 1.791 | 2.865 | 0.2517 |
| 2 | med. 2yr $\sigma_{pd}$ | 7.065 | 30.4 | **0.0756** | 15.296 | 1.586 | 2.487 | 0.2530 |
| 3 | std. ret. | 5.943 | 36.9 | 0.0636 | **12.962** | 1.313 | 1.919 | 0.2690 |
| 4 | P/D ratio | 5.529 | **46.7** | 0.0571 | 11.727 | 1.219 | 1.687 | 0.2841 |
| 5 | TIPS sp. | n.c. | n.c. | n.c. | n.c. | **0.424** | 0.146 | n.c. |
| 6 | TIPS rx. | 2.801 | 319.6 | 0.0213 | 4.616 | 0.533 | **0.822** | 0.3733 |
| 7 | pred. $\hat{\beta}_1$ | 1.742 | 43.3 | 0.0353 | 7.153 | 0.397 | 0.411 | **0.2663** |
| *Group 2: pairs anchored on TIPS 10–5yr spread* | | | | | | | | |
| 8 | eq. ret., TIPS sp. | n.c. | n.c. | n.c. | n.c. | **0.625** | 0.643 | n.c. |

*continued on next page*

Table 7 (continued)

| Run | Targeted moments | $\bar{r}_e$ | $\overline{P/D}$ | med $\sigma_{pd}^{2y}$ | $\sigma_e$ | TIPS sp. | TIPS rx. | $\hat{\beta}_1$ |
|---|---|---|---|---|---|---|---|---|
| Empirical target | | 8.384 | 53.8 | 0.0671 | 12.497 | 0.457 | 0.708 | 0.2152 |
| 9 | med. 2yr $\sigma_{pd}$, TIPS sp. | 3.254 | 88.1 | **0.0772** | 16.156 | **0.570** | 0.672 | 0.2251 |
| 10 | std. ret., TIPS sp. | 2.677 | 39.2 | 0.0649 | **13.018** | **0.504** | 0.643 | 0.2126 |
| 11 | P/D ratio, TIPS sp. | 5.343 | **48.7** | 0.0561 | 11.526 | **1.173** | 1.601 | 0.2846 |
| 12 | TIPS sp., TIPS rx. | n.c. | n.c. | n.c. | n.c. | **0.353** | **0.529** | n.c. |
| 13 | TIPS sp., pred. $\hat{\beta}_1$ | 2.804 | 39.1 | 0.0687 | 13.750 | **0.484** | 0.654 | **0.1999** |
| *Group 3: pairs anchored on mean P/D ratio* | | | | | | | | |
| 14 | eq. ret., P/D ratio | **5.676** | **46.7** | 0.0581 | 11.948 | 1.250 | 1.744 | 0.2843 |
| 15 | P/D ratio, med. 2yr $\sigma_{pd}$ | 5.640 | **46.7** | **0.0580** | 11.919 | 1.242 | 1.729 | 0.2842 |
| 16 | P/D ratio, std. ret. | 5.679 | **46.7** | 0.0581 | **11.943** | 1.252 | 1.745 | 0.2845 |
| 17 | P/D ratio, TIPS sp. | 5.343 | **48.7** | 0.0561 | 11.526 | **1.173** | 1.601 | 0.2846 |
| 18 | P/D ratio, TIPS rx. | 4.802 | **46.8** | 0.0518 | 10.616 | 1.064 | **1.413** | 0.2837 |
| 19 | P/D ratio, pred. $\hat{\beta}_1$ | 5.227 | **46.7** | 0.0556 | 11.413 | 1.150 | 1.568 | **0.2821** |
| *Group 4: triples anchored on TIPS spread + P/D ratio* | | | | | | | | |
| 20 | eq. ret., P/D ratio, TIPS sp. | **5.671** | **46.7** | 0.0581 | 11.945 | **1.249** | 1.742 | 0.2842 |
| 21 | P/D ratio, med. 2yr $\sigma_{pd}$, TIPS sp. | 5.343 | **48.7** | **0.0561** | 11.526 | **1.173** | 1.601 | 0.2846 |
| 22 | P/D ratio, std. ret., TIPS sp. | 5.626 | **46.7** | 0.0577 | **11.854** | **1.240** | 1.725 | 0.2845 |

*continued on next page*

Table 7 (continued)

| Run | Targeted moments | $\bar{r}_e$ | $\overline{P/D}$ | med $\sigma_{pd}^{2y}$ | $\sigma_e$ | TIPS sp. | TIPS rx. | $\hat{\beta}_1$ |
|---|---|---|---|---|---|---|---|---|
| Empirical target | | 8.384 | 53.8 | 0.0671 | 12.497 | 0.457 | 0.708 | 0.2152 |
| 23 | P/D ratio, TIPS sp., TIPS rx. | 4.965 | **46.8** | 0.0528 | 10.818 | **1.099** | **1.475** | 0.2834 |
| 24 | P/D ratio, TIPS sp., pred. $\hat{\beta}_1$ | 5.200 | **46.7** | 0.0556 | 11.406 | **1.143** | 1.557 | **0.2815** |

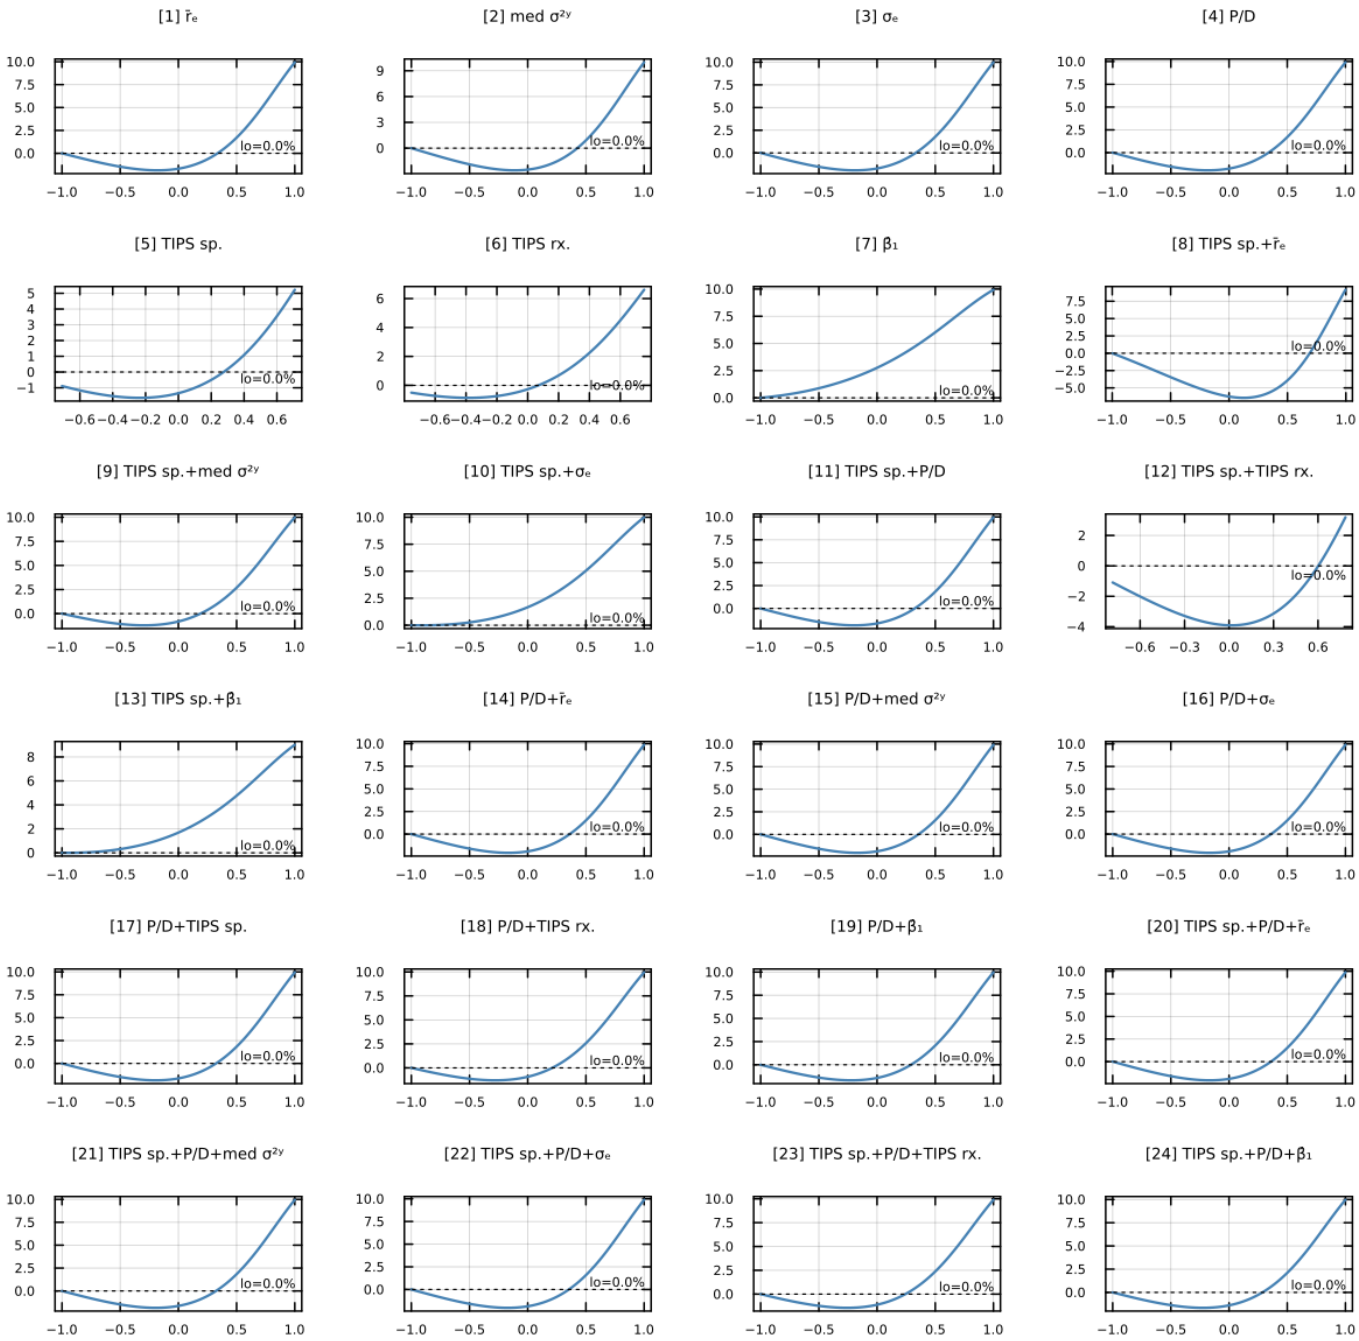


Figure 2: Instantaneous Risk-Free Rate Profile across the 24 variations, showing that the pattern of the risk-free rate is relatively common across the different parameterizations. This is also driven by the fact that all runs start from the same initial parameter vector.

# E Simulation

This section provides additional details on the simulation of the baseline model's state variable and asset prices.

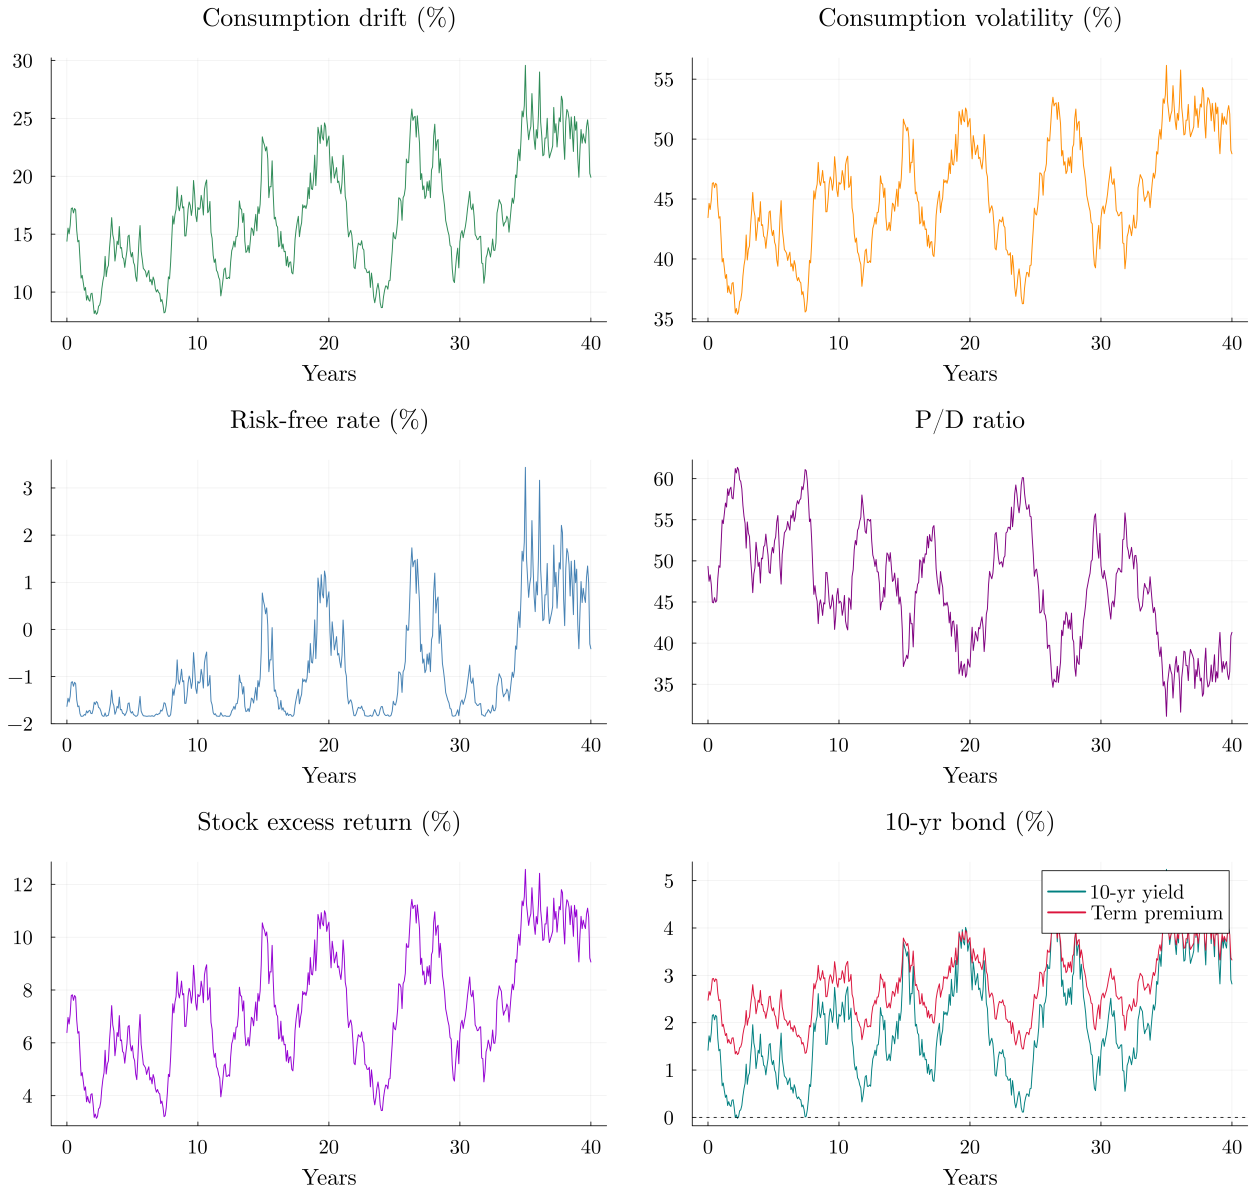


Figure 3: Variables based on a simulated path for forty years under the baseline parameterization.

# F Expectation-Hypothesis Regression Results

Tables 8 and 9 report the FB and CS expectation-hypothesis regression slope estimates at the baseline calibration, together with their cross-simulation standard deviation and the empirical target. At the baseline parametrization the sign of the estimates is consistent with the targets: FB betas are positive (though below unity) and CS betas are negative, both indicating a violation of the expectations hypothesis. Across different parametrizations, however, while the expectations hypothesis continues to be violated, the exact pattern of the estimates can vary significantly. The underlying predictability of bond returns and the state-dependent nature of bond term premia and risk premia are illustrated more reliably in the figures in the main text.

Table 8: Fama–Bliss expectation-hypothesis regressions: model vs. data. OLS of quarterly log excess bond return on forward–spot spread. Model column reports mean slope across 2,000 simulated 60-year samples; standard deviations across simulations in parentheses. $\beta = 1$ under the expectations hypothesis.

| Maturity | $\bar{\beta}$ | (SD) | Target |
|---|---|---|---|
| 2-year | 0.927 | (1.221) | 1.330 |
| 3-year | 0.889 | (1.502) | 1.710 |
| 4-year | 0.725 | (1.801) | 1.840 |
| 5-year | 0.421 | (2.187) | 1.690 |

Table 9: Campbell–Shiller expectation-hypothesis regressions: model vs. data. OLS of scaled yield change on yield spread. Model column reports mean slope across 2,000 simulated 60-year samples; standard deviations across simulations in parentheses. $\beta < 0$ indicates violation of the expectations hypothesis.

| Horizon | $\bar{\beta}$ | (SD) | Target |
|---|---|---|---|
| 3-month | -1.216 | (1.392) | -1.030 |
| 6-month | -0.851 | (1.722) | -1.160 |
| 1-year | -1.151 | (1.763) | -1.410 |
| 2-year | -0.991 | (1.902) | -1.920 |
| 4-year | -0.954 | (2.680) | -2.830 |

# G Intermediary Consumption Volatility

This appendix provides indirect evidence on the level and dynamics of the consumption volatility of the marginal intermediary investor. Following He *et al.* (2017), the

1-year rolling equity return volatility of NY Fed primary dealer holding companies is used directly as a proxy for the consumption volatility $\sigma_c$ of the marginal investor. At each date, the implied state variable $x_t$ is deduced by inverting $\sigma_c(x)$ using this empirical series, and the corresponding model-implied asset prices are then computed.

The mean intermediary equity return volatility in the sample is 23.1%, somewhat below the model's unconditional mean $\sigma_c$ of 43.9%, so the mapping places the implied $x_t$ predominantly in the lower half of the state space on average. However, during stress episodes the data exhibit considerably more extreme values: the series peaks at 72.6% near the 2008 financial crisis, consistent with the interpretation that crises correspond to large positive realizations of $x_t$.

Table 10 summarises the volatility statistics and the correlations between the model-implied and empirical asset-price series. The implied P/D ratio displays a positive correlation of 0.48 (levels) and 0.35 (changes) with its empirical counterpart, suggesting that the state variable captures a meaningful share of equity valuation dynamics. The 10–5 yr TIPS spread also displays a positive levels correlation of 0.33. The 10-year TIPS yield correlations are weaker, reflecting that the yield level is influenced by factors—such as the trend in real rates—not captured by the single-factor model. The correlations are moderate rather than large.

The model was not calibrated to match this specific intermediary series, so an imperfect fit is expected. Most importantly, the model sets $\gamma = 2$, which may not reflect the risk aversion of the specific intermediaries in He *et al.* (2017). This does not rule out the existence of other marginal investors with risk aversion close to this value. A more disciplined approach would jointly match the consumption process and the risk aversion parameter to the characteristics of the relevant intermediaries, so that the implied connection to asset prices is as direct as possible.

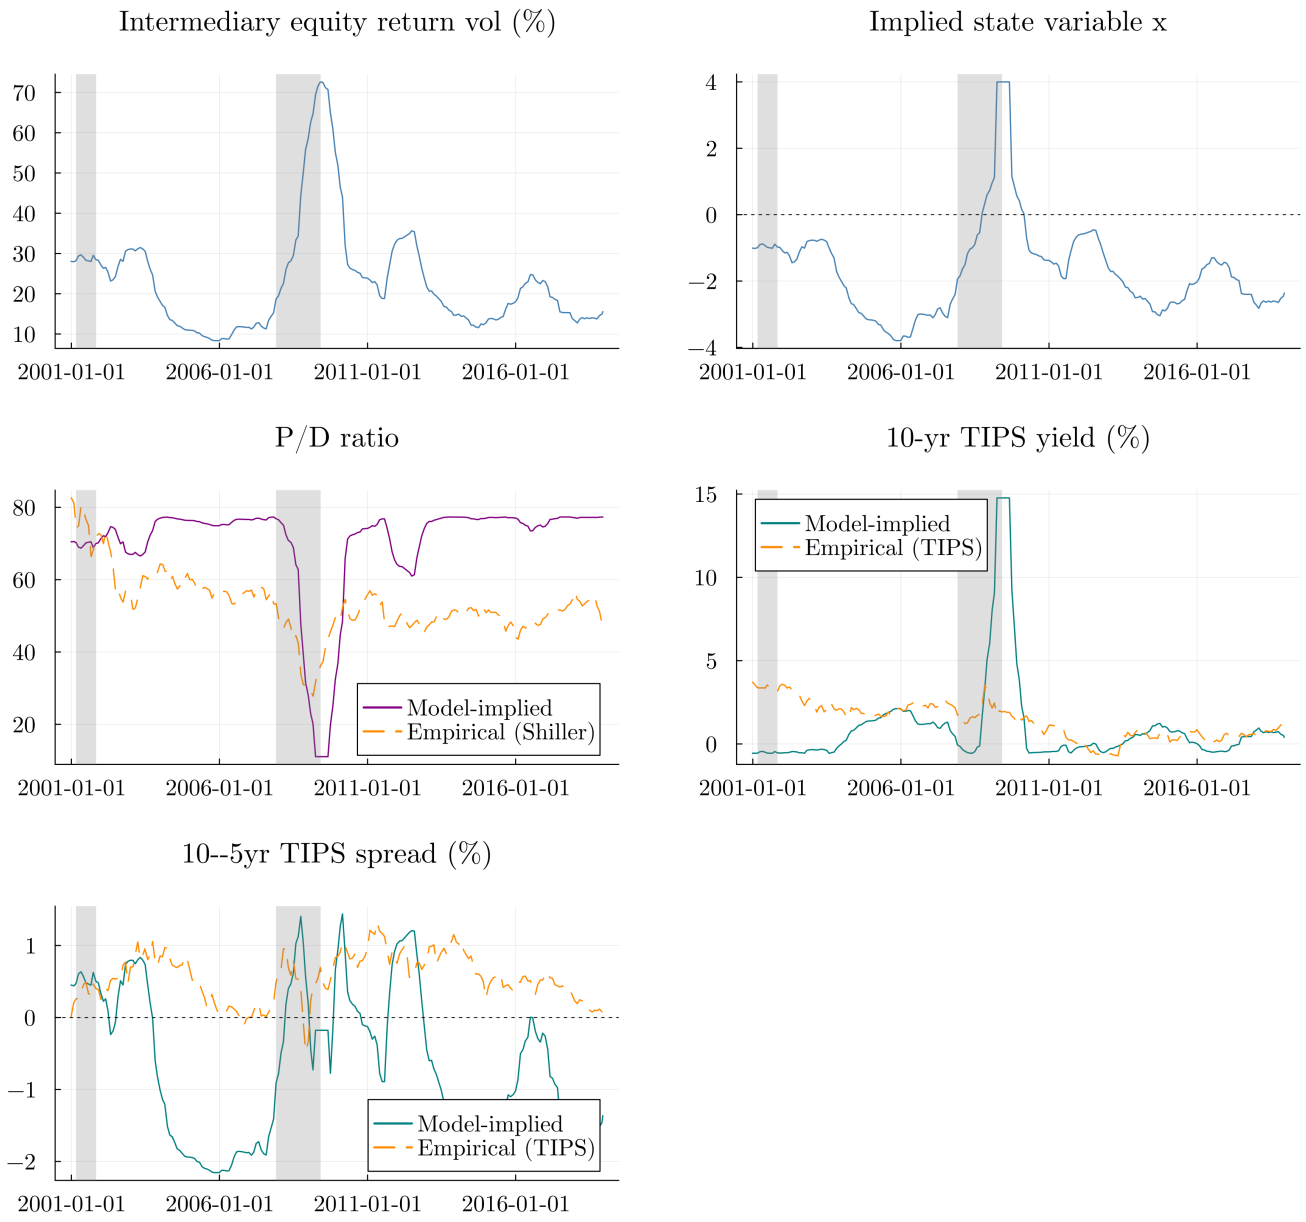


Figure 4: Intermediary-implied state variable and asset prices. *Top left:* 1-year rolling equity return volatility of NY Fed primary dealer holding companies from He *et al.* (2017). *Top right:* implied state variable $x_t$ obtained by inverting $\sigma_c(x)$ at each date using the empirical equity return volatility. On dates when the empirical volatility exceeds the model's soft-capped maximum ($\approx 67\%$), exact inversion is infeasible and $x_t$ is set to the grid ceiling $x = 4$. *Middle left/right and bottom:* model-implied P/D ratio, 10-year TIPS yield, and 10–5 yr TIPS spread (solid lines) overlaid on their empirical counterparts (dashed lines). Sample: 2001–present.

Table 10: Indirect evidence on intermediary consumption volatility. Panel A reports summary statistics for the 1-year rolling equity return volatility of NY Fed primary dealer holding companies, computed from the He *et al.* (2017) daily return series (sample: 2001–present). AR(1) coefficient estimated on the monthly subsampled series. Panel B reports correlations between model-implied and empirical asset prices. The implied series are obtained by inverting $\sigma_c(x)$ at each date using the empirical intermediary equity return volatility, then evaluating bond and stock prices at the implied state. Yield and spread data: Federal Reserve TIPS yield curve (FEDS 200805); P/D ratio: Shiller (2015).

*Panel A: Intermediary equity return volatility (%)*

| Mean | Std | Min | Max | AR(1)[a] |
|---|---|---|---|---|
| 23.1 | 13.7 | 8.3 | 72.6 | 0.890 |

*Panel B: Correlations — model-implied vs. empirical*

| Series | Levels | Changes |
|---|---|---|
| P/D ratio | 0.484 | 0.345 |
| 10-yr TIPS yield | 0.151 | -0.009 |
| 10–5yr TIPS spread | 0.332 | 0.047 |

[a]AR(1) on monthly series; reported value is $\phi^{12}$, the implied annual persistence.

# References


Abrahams, M., Adrian, T., Crump, R. K., Moench, E. and Yu, R. (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics*, **84**, 182–200.

Adrian, T., Etula, E. and Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, **69** (6), 2557–2596.

— and Shin, H. S. (2014). Procyclical leverage and value-at-risk. *The Review of Financial Studies*, **27** (2), 373–403.

Ait-Sahalia, Y., Parker, J. A. and Yogo, M. (2004). Luxury goods and the equity premium. *The Journal of Finance*, **59** (6), 2959–3004.

Backus, D. K., Gregory, A. W. and Zin, S. E. (1989). Risk premiums in the term structure: Evidence from artificial economies. *Journal of Monetary Economics*, **24** (3), 371–399.

Bansal, R. and Shaliastovich, I. (2013). A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets. *The Review of Financial Studies*, **26** (1), 1–33.

— and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, **59** (4), 1481–1509.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, **121** (3), 823–866.

Basak, S. and Cuoco, D. (1998). An equilibrium model with restricted stock market participation. *The Review of Financial Studies*, **11** (2), 309–341.

Benzoni, L., Collin-Dufresne, P. and Goldstein, R. S. (2011). Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics*, **101** (3), 552–573.

Campbell, J. Y. (2006). Household finance. *The journal of finance*, **61** (4), 1553–1604.

— and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, **107** (2), 205–251.

— and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The review of financial studies*, **1** (3), 195–228.

— and — (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, **58** (3), 495–514.

Chen, Y., Cosimano, T. F. and Himonas, A. A. (2010). Continuous time one-dimensional asset-pricing models with analytic price–dividend functions. *Economic Theory*, **42** (3), 461–503.

—, —, — and Kelly, P. (2009). Asset pricing with long run risk and stochastic differential utility: An analytic approach. *Available at SSRN 1502968*.

Cochrane, J. H. (2009). *Asset Pricing: Revised edition*. Princeton University Press.

d' Amico, S., Kim, D. H. and Wei, M. (2018). TIPS from TIPS: The informational content of treasury inflation-protected security prices. *Journal of Financial and Quantitative Analysis*, **53** (1), 395–436.

Duffie, D. and Epstein, L. G. (1992). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, pp. 353–394.

Ellison, M. and Tischbirek, A. (2021). Beauty contests and the term structure. *Journal of the European Economic Association*, **19** (4), 2234–2282.

Fama, E. F. and Bliss, R. R. (1987). The information in long-maturity forward rates. *The American Economic Review*, pp. 680–692.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, **127** (2), 645–700.

Gürkaynak, R. S., Sack, B. and Wright, J. H. (2010). The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics*, **2** (1), 70–92.

Haddad, V. and Muir, T. (2021). Do intermediaries matter for aggregate asset prices? *The Journal of Finance*, **76** (6), 2719–2761.

Haliassos, M. and Bertaut, C. C. (1995). Why do so few hold stocks? *the economic Journal*, **105** (432), 1110–1129.

He, Z., Kelly, B. and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, **126** (1), 1–35.

— and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, **103** (2), 732–770.

Kargar, M. (2021). Heterogeneous intermediary asset pricing. *Journal of Financial Economics*, **141** (2), 505–532.

Kekre, R., Lenel, M. and Mainardi, F. (2022). Monetary policy, segmentation, and the term structure.

Kraft, H., Seiferling, T. and Seifried, F. T. (2017). Optimal consumption and investment with epstein–zin recursive utility. *Finance and Stochastics*, **21**, 187–226.

Lettau, M. and Wachter, J. A. (2011). The term structures of equity and interest rates. *Journal of Financial Economics*, **101** (1), 90–113.

Mehra, R. (2007). The equity premium puzzle: A review. *Foundations and Trends® in Finance*, **2** (1), 1–81.

— and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, **15** (2), 145–161.

Melissinos, E. (2023a). A perturbation solution for the pricing of long-term bonds under recursive utility.

— (2023b). Real term premia in consumption-based models. *Available at SSRN 4582708*.

Mitra, I. and Xu, Y. (2024). A theory of the term structure of interest rates under limited household risk sharing. *The Review of Financial Studies*, **37** (8), 2461–2509.

Mogensen, P. and Riseth, A. (2018). Optim: A mathematical optimization package for julia. *Journal of Open Source Software*, **3** (24).

Pflueger, C. E. and Viceira, L. M. (2016). Return predictability in the treasury market: Real rates, inflation, and liquidity. *Handbook of Fixed-Income Securities*, pp. 191–209.

Piazzesi, M. and Schneider, M. (2006). Equilibrium yield curves. *NBER Macroeconomics Annual*, **21**, 389–442.

Rietz, T. A. (1988). The equity risk premium a solution. *Journal of Monetary Economics*, **22** (1), 117–131.

Schneider, A. (2022). Risk-sharing and the term structure of interest rates. *The Journal of Finance*, **77** (4), 2331–2374.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends?

— (2015). Irrational exuberance: Revised and expanded third edition.

Tsai, J. and Wachter, J. A. (2018). Pricing long-lived securities in dynamic endowment economies. *Journal of Economic Theory*, **177**, 848–878.

Vayanos, D. and Vila, J.-L. (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica*, **89** (1), 77–112.

Wachter, J. A. (2006). A consumption-based model of the term structure of interest rates. *Journal of Financial Economics*, **79** (2), 365–399.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of monetary economics*, **24** (3), 401–421.

Zhao, G. (2020). Ambiguity, nominal bond yields, and real bond yields. *American Economic Review: Insights*, **2** (2), 177–192.