

High Volatility is the Key:

A Joint Treatment of Asset Pricing Puzzles*

Errikos Melissinos[†]

August 2026

Abstract

Asset pricing puzzles arise when the marginal investor is identified with the average consumer. Increasing evidence indicates that the marginal investor is not the average household, which makes it natural to model this investor as an agent with much higher consumption volatility. In a mostly standard model whose essential feature is high volatility, none of the puzzles studied here arise: the equity premium, the risk-free rate, the bond premium, excess volatility, and stock return predictability. Almost all of the corresponding moments are also matched closely. The theory is general, while the calibration adopts the financial intermediary interpretation.

JEL: C65, E43, E44, G12

Keywords: term premia, bond premium puzzle, stochastic volatility, intermediary asset pricing, consumption-based asset pricing, high consumption volatility, equity premium, risk-free rate puzzle, return predictability, excess volatility

*For helpful comments I thank Han Özsoylev, Christian Schlag and Volker Wieland. An earlier version of this paper was circulated under the title “Real Term Premia in Consumption-Based Models” (SAFE Working Paper No. 413). The comparison of model variations and the perturbation solution method used there now appear in the Internet Appendix.

[†]Özyegin University
website: <https://errikos-melissinos.com>

1 Introduction

The asset pricing literature has accumulated a long list of puzzles. This article suggests that most of them disappear under a simple and natural change to one important assumption. In particular, the puzzles arise when the marginal investor is identified with the average consumer. When the marginal investor is not the average consumer, reasonable specifications exist under which they do not arise. Indeed, several independent strands of literature already conclude that the marginal investor differs from the average household. This article addresses several of the most fundamental puzzles with a mostly standard asset pricing model whose essential feature is the high and time-varying consumption volatility of the marginal investor, a feature that is reasonable under several current views about who that investor is. No puzzle arises in this model, and most of the moments connected to the puzzles are matched closely.

The relevant literature is increasingly clear that the agents pricing risky assets are not ordinary households. Regular households have limited stock market participation (Haliassos and Bertaut 1995; Campbell 2006). Wealthy households, who are more plausibly marginal in asset markets, have consumption that is both far more volatile than the aggregate and more strongly correlated with stock returns (Ait-Sahalia, Parker and Yogo 2004). 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). The intermediary asset pricing approach identifies financial intermediaries as the relevant marginal investors (He and Krishnamurthy 2013).

Before turning to the model, it helps to state the puzzles precisely. A puzzle is a pattern in central financial variables (such as stock returns, the risk-free rate, and bond yields) that standard macroeconomic models fail to replicate,

either qualitatively or by a wide margin 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, in the data long-term bonds also exhibit positive excess returns, which are large and time-varying, while most standard theoretical models predict excess returns for bonds that are negative and small in absolute value. This is the bond premium puzzle (Backus, Gregory and Zin 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; Joyce, Kaminska and Lildholdt 2012), while in the USA most evidence points to a significant, mostly positive and time-varying real term premium (Abrahams, Adrian, Crump, Moench and Yu 2016; d' Amico, Kim and Wei 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.

Notably, addressing the bond premium puzzle for real bonds 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

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

²The 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.

theoretically hedge consumption risk and command a negative risk premium instead of a positive one. As shown in Section 5, 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 sizeable 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. Up to now, two main alternatives have been used to address this challenge.³ The first relies on an inflation channel, which by definition does not apply to real bonds. The second alternative relies on highly time-varying effective risk aversion, such as in habit-based frameworks (Campbell and Cochrane 1999; Wachter 2006).⁴ 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).

This article makes two main contributions. The first is to flesh out a characteristic that is common to many successful explanations of these puzzles: the marginal investor faces high, time-varying volatility in consumption or wealth. 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, Lenel and Mainardi 2024; Schneider 2022),

³A more extensive review of this literature is provided in the Internet Appendix.

⁴It 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).

and supply-and-demand driven models (Vayanos and Vila 2021), all converge on this idea. The current article distills the common feature into a more abstract consumption-based model, in two parts: high, time-varying consumption volatility is analyzed in its own right, and when it comes to the calibration of the model, the intermediary perspective is used. The second contribution is a joint treatment of the puzzles described above. Several frameworks pursue joint treatments, including habit formation (Campbell and Cochrane 1999), long-run risks (Bansal and Yaron 2004), and variable rare disasters (Gabaix 2012; Wachter 2013). To the author’s knowledge, however, the framework of this article is the only one that jointly tackles these basic puzzles together with a positive real bond premium.

Two questions are worth separating. The first is whether the puzzles arise at all. A puzzle is taken to arise when the mismatch between model and data is glaring, as with the equity premium above, where a representative agent calibrated to aggregate consumption delivers a premium more than an order of magnitude below the one observed. Under this criterion none of the puzzles studied here arise in the model. The second question is whether the model also recovers the exact values of the moments. On this front it largely succeeds as well, and almost all of them are matched closely. The equity premium is the exception, since it still falls short of the data by a statistically significant margin.

This article uses an exogenous process for the consumption of the marginal investor, and the process is deliberately reduced-form. The theory does not require consumption volatility to be high. It is the matching of the moments that requires consumption volatility to be high and time-varying. What separates the marginal investor from the representative household is therefore a requirement imposed by the moments, not a feature built into the theory. Because the theory is silent on who that investor is, several interpretations are admissible. It could be a financial intermediary whose risk exposure and performance deteriorate together in high-volatility states. It could be a group of wealthy households whose consumption

is volatile and strongly linked to financial markets (Ait-Sahalia *et al.* 2004). It could also reflect households that cannot share idiosyncratic risk, as in Mitra and Xu (2024), where such households generate bond premia. The analysis does not depend on which interpretation is adopted, and the exogenous consumption process abstracts from the specific frictions behind each one. The intermediary interpretation is adopted at the level of the calibration, and the volatility measured from intermediary equity returns in the data is largely consistent with the magnitude the moments require.

The exogenous state variable that drives consumption volatility also drives the consumption drift. In contrast to the pro-cyclical interest rates noted above, the real rate increases with consumption volatility. The exact form of this relationship is not imposed but arises from the calibrated parameters, as described further in the article. These features generate a sizeable 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 negatively correlated with the consumption shocks of the marginal investor. A negative shock raises the volatility state, which raises the real rate and depresses bond prices. Bonds, therefore, perform well (badly) in good (bad) times for the marginal investor. 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. All of this variation traces back to a single state variable. Keeping only one state variable makes the model simpler and establishes how much explanatory power a single dimension of risk can have. It also comes at a natural cost, as one state variable cannot capture every dimension of asset pricing.

To perform the calibration, the article adopts the intermediary perspective, which supplies the data. The calibration, described in Section 3, proceeds in two steps. The persistence of the volatility state and the entire volatility schedule

are first measured from the equity returns of NY Fed primary dealers (He, Kelly and Manela 2017). These dealer data are used in the intermediary asset pricing literature, where their relevance for pricing a wide cross-section of assets has been documented by Adrian, Etula and Muir (2014), He *et al.* (2017), and Haddad and Muir (2021). Dealer equity volatility is treated as what the model says it is (the return volatility of the marginal investor’s wealth) and the model inverts it into the consumption volatility schedule. The two remaining parameters are then calibrated to eight asset pricing moments, so the system is heavily overidentified. Despite this, most of the moments are matched closely and jointly, which is itself evidence for the intermediary perspective. In addition, the state variable recovered from dealer volatility alone comoves positively with observed price–dividend ratios and TIPS spreads. The model closely matches valuations, return volatility, predictability, and the real term structure simultaneously, yet it falls short of the equity premium, even though high consumption volatility is the classic remedy for exactly that moment. In the current setup with one state variable, the equity premium level cannot be matched perfectly together with the term premia. The remaining gap either requires additional state variables or extra mechanisms (Section 4.2).

2 Model Setup

2.1 Marginal Investor

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 markets due to frictions. The identity of the marginal investor is left open in this section. Because the model is a simple consumption-based framework, it can

accommodate any of the interpretations discussed in the introduction, whether a financial intermediary, a group of wealthy households, or households unable to share idiosyncratic risk. The intermediary interpretation is adopted later in the calibration.

2.2 Utility

The lifetime utility of the agent is given by:

$$U_0 = E_0 \int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma} - 1}{1-\gamma} dt, \quad (1)$$

where C_t is the consumption flow at time t , γ is the coefficient of relative risk aversion, and ρ 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:

$$d \log(C_t) = dc_t = \mu_c(x_t)dt + \sigma_c(x_t)dW_{ct} \quad (2)$$

where c_t is log consumption, x_t is the state variable, W_{ct} is a Wiener process associated with consumption, $\mu_c(x_t)$ is the consumption drift, and $\sigma_c(x_t)$ is the consumption diffusion, abbreviated below as μ_{ct} and σ_{ct} . The specific functional forms of μ_{ct} and σ_{ct} are given in the Appendix. They depend on parameters whose calibration is described in Section 3. As a result of the calibration, both μ_{ct} and σ_{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:

$$\begin{aligned} dx_t &= \log(\phi)x_t dt + \sigma_{xt} dW_{xt} \\ dW_{xt} dW_{ct} &= \rho_{cx} dt \end{aligned} \tag{3}$$

where σ_{xt} is the state variable diffusion, W_{xt} is a Wiener process associated with the state variable, and ρ_{cx} 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 $0 < \phi < 1$.⁵

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:⁶

$$\frac{d\Lambda_t}{\Lambda_t} = \left(-\rho - \gamma\mu_{ct} + \frac{\gamma^2}{2}\sigma_{ct}^2 \right) dt - \gamma\sigma_{ct} dW_{ct} \tag{5}$$

The instantaneous risk-free rate then follows as:

$$r(x_t)dt = -E_t \left[\frac{d\Lambda_t}{\Lambda_t} \right] = \left(\rho + \gamma\mu_{ct} - \frac{\gamma^2}{2}\sigma_{ct}^2 \right) dt \tag{6}$$

⁵This is similar to an AR(1) where $x_{t+1} = \phi x_t + \epsilon_t$.

⁶The pricing follows the notation and the approach of Cochrane (2009).

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:

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

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 m .⁷ Using Itô's lemma on the bond price yields the following:

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

where subscripts on B denote partial derivatives with respect to the corresponding variable.⁸ 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):

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

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. Based on the expressions above, the instantaneous

⁷Although the article works with zero-coupon bonds, coupon-bearing bonds are priced by the same solution. A coupon bond is a portfolio of zero-coupon claims, so its price follows by summing the prices of the zero-coupon bonds corresponding to its payments, or by integrating them in the case of a continuous coupon flow.

⁸For simplicity the dependence on time is not explicitly shown, but maturity is connected to time through $\partial/\partial t = -\partial/\partial m$.

expected excess return of the bond before maturity can be derived as:

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

Here the risk premium is connected to the level of risk aversion and the covariance between the instantaneous changes in consumption and the bond price, as the right-hand side of the equation includes both of these elements. The term premium is defined as:⁹

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

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 premium:

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

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) $\rho_{cx} \gamma \sigma_{ct} \sigma_{xt}$ can be compared to $\log(\phi)x_t$, as both terms enter additively in the equation. The former is the term premium component, as it drives the term premium, while the latter term is the expectation component, as it drives the deviation of yields from the short-term rate that is due to expected future changes in that rate. If the term

⁹Based on the Feynman-Kac formula, the equation can also be written like: $TP(x_t, m) = -\frac{1}{m} \log \left[\frac{\mathbb{E}_t^* \left[\exp \left\{ -\int_0^m r(x_{t+s}) ds \right\} \right]}{\mathbb{E}_t \left[\exp \left\{ -\int_0^m r(x_{t+s}) ds \right\} \right]} \right]$, where the numerator uses the risk-neutral measure.

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 (ρ_{cx}) is zero, then the term premium is zero and the expectations hypothesis holds.¹⁰ 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. The Internet Appendix 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:

$$\begin{aligned} \frac{dD_t}{D_t} &= \mu_D dt + \sigma_D dW_{Dt} \\ dW_{Dt}dW_{ct} &= \rho_{cD}dt, \quad dW_{Dt}dW_{xt} = \rho_{xD}dt \end{aligned} \tag{13}$$

where D_t is the dividend flow at time t , μ_D is the drift of the dividend process, σ_D is the dividend diffusion, W_{Dt} is a Wiener process associated with the dividend process, ρ_{cD} is the correlation between the noise processes for consumption and the dividend, and ρ_{xD} is the correlation between the noise processes for the state variable and the

¹⁰Indeed, 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.

dividend. For simplicity, μ_D and σ_D are assumed to be constant.¹¹ 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 at the time, so $\hat{s}(x, 0) = 1$, which uniquely pins down the solution.¹² By Itô’s lemma, $\hat{s}$ follows the process:

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

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):¹³

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

where D is the dividend paid at time $t + m$. Comparing Equation (15) with

¹¹The model can still be solved in a similar way if they depend on the state variable x_t .

¹²Trying 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, Cosimano and Himonas, 2010). The approach followed here avoids this.

¹³This 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.

Equation (9) shows an analogous structure, while the differences stem from the stochastic dividend process, which adjusts the discount rate and adds a term to the drift. The differential equation can be solved in the same way as the bond pricing equation. The risk premium of the dividend strip can also be derived from the pricing equation:¹⁴

$$\begin{aligned} \mathbb{E}\left[\frac{d(\hat{s}D)}{\hat{s}D}\right] + \mathbb{E}\left[\frac{d\Lambda}{\Lambda}\right] &= \underbrace{\mathbb{E}\left[\frac{d(\hat{s}D)}{\hat{s}D}\right] - r(x)dt}_{\equiv \text{Expected Excess Return}} = -\mathbb{E}\left[\frac{d\Lambda}{\Lambda} \frac{d\hat{s}}{\hat{s}} + \frac{d\Lambda}{\Lambda} \frac{dD}{D}\right] \end{aligned} \quad (16)$$

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

As with the bond, the risk premium is connected to the level of risk aversion and the covariance of the instantaneous changes in consumption with the two sources of the strip's return. The first term on the right-hand side reflects changes in the price-dividend ratio and the second reflects changes in the dividend itself. The stock price-dividend ratio is obtained by integrating the strip price-dividend ratio across all maturities:¹⁵

$$\frac{S(x_t)}{D(x_t)} = s(x_t) = \int_0^\infty \hat{s}(x_t, m) dm \quad (17)$$

where the stock's dividend flow matches the dividends paid by the dividend strips.¹⁶ The expected excess return of the stock follows correspondingly as a weighted

¹⁴The risk premium is used equivalently as the expected excess return.

¹⁵A similar approach is used in Wachter (2006).

¹⁶In 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.

average of the strip excess returns across maturities:¹⁷

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

The pricing PDEs for the bond and the strip price–dividend ratio are of the form arising from the Feynman-Kac theorem, so each PDE links the price function to the expected discounted payoff along a risk-adjusted diffusion. Both PDEs are solved numerically using finite difference methods with zero-flux Neumann boundary conditions. As a validation check, bond prices are also computed by Monte Carlo simulation of the respective SDE and the two methods are verified to agree.

2.7 The Link between the Consumption and the Wealth Process

The consumption process can be linked to the wealth process. The derivations in this subsection are used later in the calibration, which measures the volatility process from the equity returns of NY Fed primary dealers. Indeed, the consumption process could be the result of a wealth-management problem. The net worth of the investor follows a stochastic process:

$$dN_t = N_t (\mu_N(x_t) dt + \sigma_N(x_t) dW_{Nt}) - C_t dt \quad (19)$$

$$dW_{Nt} dW_{xt} = \rho_{Nx} dt$$

where μ_N and σ_N are the drift and diffusion of the return on net worth, ρ_{Nx} is the correlation between wealth shocks and state shocks, and the state variable x_t

¹⁷Alternatively 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 $\mathbb{E}[d(\Lambda S)] + \Lambda D dt = 0$, 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.

summarizes the investment opportunity set. For a levered financial intermediary with assets A_t and liabilities L_t , net worth is the equity of the institution, $N_t = A_t - L_t$ and $\sigma_N = (A_t/N_t) \sigma_A$, so net-worth volatility is directly observable as the return volatility of the institution's publicly traded equity. Because CRRA preferences are homothetic and the opportunity set depends only on x_t , the optimal consumption rule is linear in wealth, $C_t = \kappa(x_t) N_t$ (Merton 1969).¹⁸ Applying Itô's lemma to $\log C_t = \log \kappa(x_t) + \log N_t$ then delivers log consumption dynamics of exactly the form (2):

$$d \log C_t = \mu_c^N(x_t) dt + \underbrace{\sigma_N(x_t) dW_{Nt} + \frac{\kappa'(x_t)}{\kappa(x_t)} \sigma_{xt} dW_{xt}}_{= \sigma_{ct} dW_{ct}} \quad (20)$$

with a drift that depends on the state and a single composite diffusion term.¹⁹ Matching the diffusion of (20) with that of (2) delivers the translation between the two processes:

$$\sigma_{ct}^2 = \sigma_N(x_t)^2 + 2 \frac{\kappa'(x_t)}{\kappa(x_t)} \rho_{Nx} \sigma_N(x_t) \sigma_{xt} + \left(\frac{\kappa'(x_t)}{\kappa(x_t)} \right)^2 \sigma_{xt}^2 \quad (21)$$

$$\rho_{cx} \sigma_{ct} = \rho_{Nx} \sigma_N(x_t) + \frac{\kappa'(x_t)}{\kappa(x_t)} \sigma_{xt} \quad (22)$$

Under log utility the consumption–wealth ratio is constant, the terms in κ'/κ vanish, and the translation is exact state by state, $\sigma_c = \sigma_N$ and $\rho_{cx} = \rho_{Nx}$. Away from log utility the κ'/κ terms reflect intertemporal hedging through the consumption–wealth ratio, and because the model determines κ , they can be computed.

The consumption–wealth ratio κ that enters the translation (21)–(22) is determined inside the model. In equilibrium, wealth equals the value of the claim

¹⁸The value function has the form $V(N, x) = f(x) N^{1-\gamma}/(1-\gamma)$, so $\kappa = f^{-1/\gamma}$ depends only on the state. Under log utility, $\kappa = \rho$, a constant.

¹⁹The drift μ_c^N plays no further role. The consumption drift is pinned by the CRRA identity of the calibration, so matching the drift terms determines the wealth drift μ_N rather than restricting consumption.

to the investor's consumption stream, and this claim is priced with the same strip machinery as the stock. A consumption strip pays the consumption C_{t+m} at horizon m , so it is a dividend strip with consumption in the place of the dividend. Its price-consumption ratio $\hat{v}(x_t, m)$ therefore satisfies the pricing equation $E[d(\Lambda \hat{v}C)] = 0$ with terminal condition $\hat{v}(x, 0) = 1$, which is Equation (15) with μ_D , σ_D , ρ_{cD} , and ρ_{xD} replaced by $\mu_{ct} + \sigma_{ct}^2/2$, σ_{ct} , 1, and ρ_{cx} .²⁰ Aggregating the strips across horizons, exactly as in Equation (17) for the stock, gives the wealth claim:

$$N_t = E_t \int_0^\infty \frac{\Lambda_{t+m}}{\Lambda_t} C_{t+m} dm = C_t \int_0^\infty \hat{v}(x_t, m) dm = C_t v(x_t), \quad \kappa(x) = \frac{1}{v(x)}, \quad (23)$$

where $v(x)$ is the price-consumption ratio of the wealth claim, the analogue of the stock's price-dividend ratio $s(x)$. This is how the model delivers the consumption-wealth ratio $\kappa(x) = 1/v(x)$ that enters the translation (21)–(22). Section 3 measures σ_N from dealer equity returns and inverts the translation to recover the consumption volatility schedule.

3 Calibration

The parameters of the model fall into three groups. The first group is set directly. A second group, the persistence of the volatility state and the entire consumption-volatility schedule $\sigma_c(x)$, is measured from the volatility process of the marginal investor. The two remaining parameters are targeted in order to fit eight asset pricing moments, so the calibration is heavily overidentified. In general, while alternative approaches could be followed, this article draws throughout on the intermediary asset pricing literature for the calibration process. In particular,

²⁰Equivalently, solving the pricing equation forward gives $\hat{v}(x_t, m) = E_t[e^{-\rho m}(C_{t+m}/C_t)^{1-\gamma}]$. The current consumption level factors out, so only consumption growth enters the valuation and $\hat{v}$ is a function of the state and the horizon alone.

marginal investors are identified with financial intermediaries, whose balance sheets are the source of the marginal investor’s wealth risk (He and Krishnamurthy 2013; Adrian *et al.* 2014; Haddad and Muir 2021). Time is measured in years throughout.

Directly set parameters. The consumption–state correlation is set to $\rho_{cx} = -1$. Under the intermediary interpretation the marginal investor’s wealth risk is state risk, driven by a single shock, so wealth and the state are perfectly correlated, and since wealth equals the value of the claim to the consumption stream, consumption tracks wealth and shares that correlation. The interpretation therefore dictates the magnitude of ρ_{cx} but not its sign, which the article takes to be negative. This choice follows the intermediary evidence of He *et al.* (2017): a deterioration of the intermediary’s balance sheet is itself a high-volatility state, since the same shock moves both. Because the intermediary is leveraged, such a deterioration lowers the value of its wealth, so the wealth and state shocks move oppositely, $dW_{Nt} = -dW_{xt}$ (that is, $\rho_{Nx} = -1$), and consumption, financed out of that balance sheet, falls with it, giving $\rho_{cx} = -1$. Economically, $\rho_{cx} < 0$ means that positive shocks to the state, which raise consumption volatility, coincide with low consumption, so high-volatility states are high marginal-utility states for the investor. The sign of this correlation depends on the marginal investor’s leverage profile. Under the book-leverage process of Adrian *et al.* (2014), as discussed in the appendix of He *et al.* (2017), the wealth–state correlation would instead be positive, which the model would also accommodate without modification. In the appendix these choices are checked directly: the positive sign yields a clearly worse fit (Appendix D.6), and the perfect correlation is consistent with the dealer data (Appendix C).

The relative risk aversion of the marginal investor is set to $\gamma = 2$, a moderate value by asset-pricing standards, so the model does not rely on high risk aversion. Because the real-rate schedule is parameterized directly, the consumption drift is

implied by it through the CRRA identity and no longer affects prices, so preferences reach the pricing of bonds and equities only through the product $\gamma\rho_{cx}$, given the zero dividend correlations set below.²¹ With $\rho_{cx} = -1$, this product equals $-\gamma$. The rate of time preference is set to $\rho = 0.01$. 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 dividend process is also set from data. Its drift is $\mu_D = 2.89\%$, the average growth of real dividends in Shiller (2015) over 1990–2024, and its diffusion is $\sigma_D = 7.2\%$, the standard deviation of real dividend growth across calendar years in the same sample. The underlying calculations are reported in Appendix A. The correlations between dividend shocks and both consumption shocks and state-variable shocks are set to zero ($\rho_{cD} = \rho_{xD} = 0$), because the measured correlation between dividend growth and the marginal investor’s returns is close to zero (Appendix G), and shutting the channel down lets the consumption-volatility mechanism operate in isolation.

Taking the intermediary interpretation to data raises a question that the calibration has to confront. It is not clear what the consumption of a financial intermediary is. Dividends, payouts, and compensation each capture only a part of it, and an institution arguably has no consumption flow in the household sense at all. The calibration does not need to answer this question, because the model links consumption to wealth. The net worth of a publicly traded intermediary is priced continuously in the equity market, so the wealth side is observable (Section 2.7), and the model supplies the conversion between the two processes. The volatility schedule measured below should therefore be read as the consumption volatility implied by the observed wealth volatility. While financial intermediaries may not technically have a consumption process, the model-derived process can be

²¹Were ρ_{cx} a free parameter, any γ would deliver the same prices and the same fit once ρ_{cx} is rescaled to hold $\gamma\rho_{cx}$ fixed, so asset prices identify only the product. The perfect wealth–state correlation pins $\rho_{cx} = -1$ and makes γ the sole dial of the product.

understood as a consumption equivalent, analogously to the consumption of regular households.

Measured from volatility data. The persistence and the volatility schedule are measured from the daily value-weighted dealer equity returns of He *et al.* (2017), whose one-year rolling volatility serves as the observable counterpart of the wealth volatility σ_N (Section 2.7 and Appendix C). The series is an imperfect proxy, since the equity returns of one group of dealers need not track the wealth of the marginal investor exactly, and the measured schedule inherits this measurement error. Over 2000–2018 this volatility series has mean 23.1%, standard deviation 13.7%, and maximum 72.6% (reached in 2008). Its monthly AR(1) coefficient implies an annual persistence $\phi = 0.890$. Because $\sigma_N(x)$ is a monotone transform of the state over the relevant range, the persistence of the volatility series identifies the persistence of x directly. The schedule $\sigma_c(x)$ is a capped logistic function of the state (Appendix B), but its parameters are not fitted to the dealer-volatility distribution directly. Dealer volatility measures $\sigma_N(x)$, the product of $\sigma_c(x)$ and a wedge computed inside the model from the consumption–wealth ratio $\kappa(x)$ introduced in Section 2.7. The schedule parameters are instead chosen so that the stationary distribution of the model-implied $\sigma_N(x)$ matches the empirical 5th, 50th, and 95th percentiles and the mean of dealer volatility. The fit matches all four statistics within one percentage point. Because the wedge depends on the full calibration and the calibration depends on the schedule, measurement and calibration are solved as a joint fixed point, which converges within two iterations (Appendix C). The cap is imposed in the space where the data live: $\max_x \sigma_N = 0.75$, just above the sample maximum. The implied cap on consumption volatility is $\sigma_{\max} = 0.62$. The wedge runs from 1.4 to 1.8 over the central range of the state, so the measured consumption volatility is well below dealer volatility. Its stationary distribution has 5th percentile 6.1%, median 10.5%, 95th percentile 33.4%, and mean 13.8%.

Figure 1 displays the mapping: the wealth-volatility schedule $\sigma_N(x)$ against the recovered consumption-volatility schedule $\sigma_c(x)$, whose vertical gap is the wedge, and the implied consumption-wealth ratio $\kappa(x)$.




Figure 1: The wealth mapping at the baseline calibration. Left: the wealth-volatility schedule $\sigma_N(x)$, whose stationary distribution is matched to dealer equity volatility, and the consumption-volatility schedule $\sigma_c(x)$ recovered by inverting the translation (21)–(22); the wedge σ_N/σ_c exceeds one throughout and is largest near the middle of the stationary distribution, and σ_N is hump-shaped, peaking at the imposed cap of 0.75, so the inversion uses the lower branch. Right: the implied consumption-wealth ratio $\kappa(x) = 1/v(x)$. Dashed vertical lines mark $\pm$ one stationary standard deviation of the state.

Derived from asset-price targets. The real-rate schedule $r(x)$ is parameterized directly, as a logistic sharing the location and width of the volatility transition (Appendix B), and the consumption drift $\mu_c(x)$ is then implied by the CRRA Euler equation. There are two free parameters, the minimum and maximum levels of the real rate $r_{\min}$ and $r_{\max}$. They are chosen to minimize the variance-weighted distance between eight simulated and empirical moments: the equity excess return, the mean and the volatility of the P/D ratio, the volatility of returns, the mean real risk-free rate, the one-year return-predictability coefficient (all from Shiller (2015), 1990–2024), and the TIPS 10–5yr yield spread and excess holding-period return (from the zero-coupon real yield curve of Gürkaynak, Sack and Wright

(2010), 2004–2024). These eight moments are the empirical counterparts of the puzzles the model sets out to address jointly: the equity excess return speaks to the equity-premium puzzle, the level and volatility of the price–dividend ratio together with the return volatility to valuations and the excess-volatility puzzle, the predictability coefficient to return predictability, the mean real rate to the risk-free-rate puzzle, and the two TIPS moments to the positive real term premium. Each moment is weighted by the inverse of its sampling variance, estimated from model simulations. With eight moments and two parameters, the minimized objective (27.1) is reported as a summary of fit rather than driven to zero. The optimum is fully interior. Both $r_{\min}$ and $r_{\max}$ lie strictly inside their bounds, and every optimizer run converges to the same values (reported in Table 1 and, individually, in Appendix D).

One of the eight targets deserves a qualification. The model’s $r(x)$ is the marginal investor’s own shadow real rate, while the empirical counterpart used here is the one-month Treasury bill rate less inflation, and the two need not be identified with one another. The observed bill rate embeds a convenience yield and a funding cost that are not part of the real rate in the model, both of which place the investor’s shadow rate above it. Because this target enters the objective through its first moment alone, granting such a wedge amounts to raising a single number and repeating the calibration, and doing so leaves the results essentially unchanged. The calibration is therefore not an artifact of identifying the bill rate with $r(x)$. Appendix D.5 reports this exercise.

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

Table 1: Baseline parameters. The set-directly panel collects the correlation, dictated by the wealth foundation of Section 2.7, and the time-preference rate, which is pricing-inert. The measured panel comes from the volatility of NY Fed primary-dealer equity returns (He *et al.*, 2017; 2000–2018): ϕ from the AR(1) of one-year rolling volatility, and the $\sigma_c(x)$ schedule by matching the stationary distribution of the model-implied wealth volatility $\sigma_N(x)$ to the empirical volatility quantiles, with the cap $\max_x \sigma_N = 0.75$ imposed just above the sample maximum (the consumption-volatility cap $\sigma_{\max}$ is solved from it). The lower panel is calibrated to eight asset-price moments with inverse-variance weights (two free parameters).

Parameter	Symbol	Value
Set directly		
Consumption–state correlation	ρ_{cx}	−1
Time preference	ρ	0.010
Measured (via the wealth mapping)		
State-variable persistence	ϕ	0.890
Consumption vol. (floor)	$\sigma_{\min}$	0.056
Consumption vol. (cap)	$\sigma_{\max}$	0.623
Volatility transition (center)	m	0.691
Volatility transition (width)	w	0.293
Calibrated (8 moments, 2 parameters)		
Real rate (low volatility)	$r_{\min}$	0.006
Real rate (high volatility)	$r_{\max}$	0.074

from three 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 mean real risk-free rate combines the Fama–French short rate with CPI inflation. All eight moments are targeted jointly by the overidentified calibration of Section 3 (two free parameters). Figure 2 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.




Figure 2: 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 , and 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.

Moment	Target	Model (SE)
Equity excess return ($\bar{r}_e$, %) ²²	8.384	3.021 (1.601)
Mean P/D ratio ($\bar{P}/\bar{D}$)	53.85	60.56 (8.57)
Median 2-yr std log(P/D) (med σ_{pd}^{2y})	0.0671	0.0603 (0.0144)
Std. stock return (σ_e , %)	12.50	14.86 (2.00)
Mean real risk-free rate ($\bar{r}_f$, %) ²³	0.030	1.620 (0.629)
TIPS 10–5yr yield spread (TIPS sp., %)	0.457	0.422 (0.037)
TIPS 10/5yr excess return (TIPS rx., %)	0.708	0.754 (0.237)
Return pred. slope ($\hat{\beta}_1$)	0.215	0.250 (0.142)

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. The P/D ratio and $\hat{\beta}_1$ are dimensionless. The Model column reports the mean across 2,000 independent simulated paths of $T = 35$ years. SE is the cross-simulation standard deviation of the moment estimator, measuring sampling variability in a 35-year window.

4.1 Consumption dynamics and the risk-free rate

Consumption volatility $\sigma_c(x)$ increases monotonically with the state variable. What reproduces the measured distribution of dealer-return volatility by construction is the implied wealth volatility $\sigma_N(x)$. Consumption volatility itself, recovered by inverting the wealth mapping, runs from 6.1% at the 5th percentile through 10.5% at the median to 33.4% at the 95th, with an unconditional mean of 13.8%, well below dealer-return volatility, but still roughly an order of magnitude above 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 $\mu_c(x)$, implied by the rate schedule, increases monotonically from small positive values in low-volatility states to strongly positive values at high x .

The risk-free rate stays within two-thirds of a percentage point of its floor of +0.6% over the entire lower half of the stationary distribution, passes through

+1.2% at the stationary median, and rises materially only in the high-volatility tail: +3.9% at the 95th percentile of the stationary distribution and +5.5% at the 99th. The model therefore does not predict rising real rates in ordinary times. The rate rise is confined to the most volatile states, as discussed in Section 4.4. The model also does not generate a risk-free-rate puzzle. The mean real rate is 1.6% (above the near-zero target, as discussed above, but far from the double-digit rates that standard calibrations imply) and its time-series variation is tame, as illustrated by the sample simulation in Appendix E. The remaining gap is one of level, and the observed bill rate is not the right yardstick for it, since it omits the convenience yield that a leveraged marginal investor does not earn. Appendix D.5 re-runs the calibration with the target raised by a plausible wedge, which narrows the gap while leaving the fit of the other moments essentially unchanged.

4.2 Equity premium, excess volatility, and return predictability

The model implies a mean P/D ratio of 60.6, about one cross-simulation standard deviation above the empirical target of 53.9. The leave-one-out analysis in Appendix D.2 shows this value is stable to the choice of the remaining targets.

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 correspondingly rises with x over the bulk of the stationary distribution (low when the state variable is low, increasing as consumption volatility and the associated risk compensation increase) and peaks in high-volatility states, where the compression of valuations begins to reduce the return exposure that earns the premium. The model falls short of the empirical target for the equity premium, delivering 3.0% against 8.4% in the data. Although this difference is statistically significant, the predicted premium is

still sizeable and far larger than a standard representative-agent calibration delivers at comparable risk aversion. The gap is not large enough to justify a puzzle. The model places the premium an order of magnitude above the standard benchmark and then falls short of the data by a factor of under three, which is a quantitative shortfall rather than an order-of-magnitude mismatch. Furthermore, many drivers of equity premia have been suggested, and some of them, such as rare disaster risk (Barro, 2006; Rietz, 1988), could be operating at the same time as the mechanism described in this article. One can nonetheless ask whether a higher premium is attainable strictly within the model, and what specifically produces the shortfall in this main calibration. When the equity premium is the sole calibration target under the measured volatility schedule, the model delivers about 4%, above the baseline but still short of the data. The full empirical value is reached only when the volatility schedule is no longer required to match the dealer data (Appendix G). The shortfall therefore reflects the measured level of consumption volatility, high but not high enough to generate the entire premium on its own, together with the joint fit of the valuation moments. Further research should explore how the gap could be closed in a richer model with additional state variables.

The model naturally produces excess volatility in stock prices relative to dividends, in the sense that the variance of stock prices exceeds what could be justified by the present value of future dividends computed with a constant discount rate, as 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 14.9%, close to the 12.5% observed in the 1990–2024 sample, and also closely matches the median 2-year rolling standard deviation of $\log(P/D)$ (0.060 in the model against an empirical value of 0.067; Table 2).²⁴

²⁴The full-sample standard deviation of $\log(P/D)$ in the data (0.27 over 1990–2024, up to 0.48

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 0.25, 0.70, and 0.90 at horizons of one, four, and eight years, capturing the well-documented pattern of predictability increasing with horizon. The levels are also consistent with the 1990–2024 estimates, although both the model-implied and the empirical slopes carry relatively large standard errors, so the comparison is not sharp.

Table 3: Return predictability regressions: OLS of the cumulative log real stock return $r_{t \rightarrow t+T}$ on the log dividend-price ratio $\log(D_t/P_t)$ at horizons $T = 1, 4, 8$ years, $r_{t \rightarrow t+T} = \alpha_T + \beta_T \log(D_t/P_t) + \varepsilon_{t+T}$. A positive slope β_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.

Horizon	Model			1990–2024			1926–2024		
	β	(s.e.)	R^2	β	(NW s.e.)	R^2	β	(NW s.e.)	R^2
1 year	0.249	(0.153)	0.101	0.215	(0.071)	0.134	0.061	(0.037)	0.023
4 years	0.704	(0.429)	0.294	0.831	(0.230)	0.454	0.209	(0.096)	0.079
8 years	0.896	(0.956)	0.375	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 greater for real term premia. Long-run risk models, for example, produce a high equity premium but fail to deliver a significant real term premium (Section 5).

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.

In the present model, long-term bonds are genuinely 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 may appear to conflict with the intuition that crises coincide with low real interest rates, a tension taken up in Section 4.4. The rate and volatility schedules share the same transition in the state, and this restriction is checked in Appendix D.3.

The yield decomposition shows that the 10-year real yield and the term premium both increase with x . At low values of x , the term premium is near zero and the yield is low. 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, with the risk-neutral yield contributing somewhat more than the term premium 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. As a further check, adjusting the yields with the liquidity-premium estimates of d’Amico *et al.* (2018) moves both targets modestly higher rather than lower (Table 6). 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. Both directly targeted moments are matched closely. 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 delivers it at the empirical magnitude. The yield decomposition into the risk-neutral yield (RNY , defined in Equation 12) and the term premium (TP , defined in Equation 11) is consistent with the pattern documented for real yields in Abrahams *et al.* (2016): a low risk-neutral yield and a considerable, state-dependent term premium, whose share of the total real yield rises from near zero in the calmest states to roughly 45% in the most volatile ones.

Appendix F reports bond return predictability regressions according to Fama and Bliss (1987) (FB) and Campbell and Shiller (1991) (CS). The model violates the expectations hypothesis in the correct direction. FB slopes are positive and above one at maturities of two to five years, overshooting the empirical values by roughly one slope unit, and CS slopes are negative and steeper than the data at every horizon. The estimates carry large cross-simulation standard deviations, and the empirical targets lie within roughly one standard deviation of the model values. These regressions should be interpreted with caution, since the empirical targets were estimated for nominal bonds and the model-implied slopes are not robust to moderate variations in the calibration.²⁵ A more reliable indicator of the model’s performance is the consistency of the yield curve and term premium functional forms with the decomposition of Abrahams *et al.* (2016), as discussed in the preceding paragraph.

²⁵The econometric literature documents that tests based on these regressions are subject to severe small-sample distortions (Bekaert, Hodrick and Marshall 1997; Bauer and Hamilton 2018). The large cross-simulation dispersion observed here suggests that the unreliability goes beyond this standard critique. Even when the model is the true data-generating process, samples of typical length produce slope estimates too dispersed to discriminate between models with very different term premium dynamics, so the regressions may have little power as a test of the expectations hypothesis. A systematic investigation of this conjecture is left for future research.

Overall, similar to He *et al.* (2017), high consumption volatility coincides with falling consumption (implying $\rho_{cx} < 0$), and this 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), still matches some of the moments, but with a single state variable its performance is limited relative to the baseline adopted here (Appendix D.6). This does not mean that the Adrian and Shin (2014) perspective is necessarily 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.

4.4 The real rate process

The calibrated schedule implies a real rate that is close to constant over ordinary states and rises only where consumption volatility is extreme. As documented above, $r(x)$ is flat near +0.6% over the calm half of the stationary distribution and rises only through the high-volatility tail, reaching +3.9% at the 95th percentile. In simulation, the short rate stays within 100 basis points of its calm floor in 66% of months and within 200 basis points in 85%. Conditional on the calm half of the state distribution its mean is +0.9% (standard deviation 0.2%), against +2.3% (standard deviation 1.1%) in the high-volatility half.

This may seem puzzling, since real rates are usually thought to fall in crises, which are precisely the states in which volatility is high. The two readings are not necessarily in conflict, because an observed rate need not measure the rate that the

marginal investor faces. Observed rates contain a convenience yield component, and that component may itself be larger in high-volatility states, when safe and liquid assets are most valued (Krishnamurthy and Vissing-Jorgensen 2012; Nagel 2016). A convenience yield that widens with volatility depresses observed rates in exactly the states where the model implies the investor’s own rate is rising, and could therefore conceal that rise. This is only one possible channel, and assessing it would require taking a stand on how convenience and liquidity premia are determined. Establishing the precise relationship between observed real rates and the rate relevant for the marginal investor is therefore left to further research.

4.5 Asset prices from the recovered state

The analysis so far has focused on matching moments. This subsection briefly explores a different question, namely how well the model captures the state of the economy over time. The calibration measures the volatility schedule from the volatility of dealer equity returns. It uses no valuation ratio or yield, and in particular none of the price series the recovered state is compared with below, which makes it possible to run the mapping in reverse. Inverting the dealer-volatility series through the wealth mapping $\sigma_N(x)$ recovers the state variable x_t at each date, and the model-implied prices that follow can be compared with their empirical counterparts. The recovered state comoves positively with the data. The model-implied price–dividend ratio has a levels correlation of 0.32 with the observed ratio (0.40 in changes), and the model-implied 10–5yr TIPS spread a levels correlation of 0.47 (0.15 in changes). A single measured state therefore carries common variation in both equity valuations and the real term structure, though the correlations are well short of unity, which is expected for a one-factor model whose state is constructed without reference to prices. Appendix H reports the full comparison. Capturing the state of the economy more accurately, for instance with additional

state variables or a richer measurement, is a natural direction for further research.

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 by itself generate positive real term premia, despite its success in explaining the equity premium.²⁶ 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:

$$V_0 = E_0 \int_0^\infty F(C_t, V_t) dt \quad (24)$$

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:

$$F(C_t, V_t) = \frac{\rho(1-\gamma)V_t}{1-1/\psi} \left(\left(\frac{C_t}{((1-\gamma)V_t)^{1/(1-\gamma)}} \right)^{1-1/\psi} - 1 \right) \quad (25)$$

where ψ 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.

²⁶As has been mentioned before, the Internet Appendix goes through many variations, and the only ones that deliver high and positive bond term premia are those with high consumption volatility.

As shown by Tsai and Wachter (2018), the value function can be expressed as:²⁷

$$V_t = \frac{C_t^{1-\gamma} e^{(1-\gamma)K(x_t)}}{1-\gamma} \quad (26)$$

where V_t increases with K , a function of x_t that captures the full dependence of the value function on the state variable.²⁸ Given the expression for the value function, Itô's lemma can be applied to derive the SDF process. Following Chen, Cosimano, Himonas and Kelly (2009), the fundamental relationship is:

$$\frac{d\Lambda}{\Lambda} = F_V(C_t, V_t)dt + \frac{dF_C(C_t, V_t)}{F_C(C_t, V_t)} \quad (27)$$

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 applying Itô's lemma to the derivative of flow utility with respect to consumption:²⁹

$$\begin{aligned} \frac{d\Lambda}{\Lambda} = & \left(\frac{\rho(-(1-\gamma\psi)e^{\frac{(1-\psi)K(x_t)}{\psi}} - \gamma\psi + \psi)}{1-\psi} - \gamma\mu_{ct} + \frac{\gamma^2\sigma_{ct}^2}{2} + \frac{\gamma(\gamma\psi-1)\rho_{cx}\sigma_{xt}\sigma_{ct}K'(x_t)}{\psi} \right. \\ & + \frac{(\gamma\psi-1)(-2\psi\log(\phi)x_tK'(x_t) + \sigma_{xt}^2((\gamma\psi-1)K'(x_t)^2 - \psi K''(x_t)))}{2\psi^2} \Big) dt \\ & - \frac{(\gamma\psi-1)\sigma_{xt}K'(x_t)}{\psi} dW_{xt} - \gamma\sigma_{ct}dW_{ct} \end{aligned} \quad (28)$$

When $\gamma = 1/\psi$ (time-separable utility), Equation (28) reduces to Equation (5). Moreover, the consumption risk term $-\gamma\sigma_{ct}dW_{ct}$ is identical to the time-separable case. Recursive utility adds the state-variable term $-\frac{(\gamma\psi-1)\sigma_{xt}K'(x_t)}{\psi}dW_{xt}$ arising from the explicit dependence of the SDF on x_t .

²⁷Similar results are common in the literature, see for example Benzoni, Collin-Dufresne and Goldstein 2011; Kraft, Seiferling and Seifried 2017.

²⁸The model is not solved fully in this section because this is not required for the analysis. A full solution of the model, as well as a perturbation-based solution method for $K(\cdot)$, is developed in the Internet Appendix.

²⁹This is done by substituting Equation (26) into the derivative of flow utility and applying Itô's lemma with consumption and the state variable as independent variables.

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:

$$\underbrace{-\frac{B_m}{B} + \log(\phi)x_t \frac{B_x}{B} + \frac{\sigma_{xt}^2}{2} \frac{B_{xx}}{B} - r(x_t)}_{\text{derived from: } E[\frac{dB}{B}] - rdt} = \left(\rho_{cx} \gamma \sigma_{ct} \sigma_{xt} + \frac{(\gamma\psi - 1)\sigma_{xt}^2 K'(x_t)}{\psi} \right) \frac{B_x}{B} \quad (29)$$

The second term in parentheses is unique to recursive utility. Furthermore, this term is in general negative. Whether the state variable drives the consumption drift or the consumption diffusion, the following qualitative relationships typically hold in equilibrium:

$$\left. \begin{aligned} (x \uparrow \Rightarrow \mu_c \uparrow) &\Rightarrow (K'(x) > 0) \text{ and } (r'(x) > 0 \Rightarrow B_x < 0) \\ (x \uparrow \Rightarrow \sigma_c \uparrow) &\Rightarrow (K'(x) < 0) \text{ and } (r'(x) < 0 \Rightarrow B_x > 0) \end{aligned} \right\} \Rightarrow K'(x)B_x < 0 \quad (30)$$

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 (29) therefore acts as a “hedging” term. This makes bonds an attractive asset for an investor worried about future states, thereby pushing the real risk premium downward.³⁰

³⁰This result is reinforced by solving the model using several calibrations in the Internet Appendix. 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 relationships could in theory be arbitrarily non-linear. In these cases, it

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 Conclusion

This article has shown that a parsimonious single-factor model, in which the marginal investor faces high and time-varying consumption volatility — measured from intermediary equity returns in the spirit of He *et al.* (2017), rather than calibrated — 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. None of these puzzles arise in the model, though that does not mean every empirical moment is precisely matched. The model reproduces the qualitative behavior that standard frameworks cannot, and for most moments it comes close to the exact values as well, despite its simplicity. For instance, it generates a positive and sizeable real term premium where those frameworks predict a near-zero or negative one. Even recursive, long-run-risk preferences, which can match the equity premium, do not by themselves generate positive real bond premia (Section 5), so the high and time-varying consumption volatility of the marginal investor is not merely one possible mechanism but an essential ingredient within this type of model.

In retrospect, this solution may have been in plain sight 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

becomes more difficult to unambiguously determine the sign of $K'(x)B_x$.

can be taken as a hint that the consumption volatility of the marginal investor is genuinely high and time-varying. In light of this reading, it becomes unsurprising that so many different strands of literature arrive at a similar conclusion, implying that the median household is not the same as the marginal investor. Ironically, this interpretation has profound consequences for regular households. In particular, their consumption is relatively stable, and the high risk premia of financial markets represent returns that far exceed what their own risk profile would demand. This suggests an extra reason why limited participation in financial markets by regular households (or other institutions not exposed to high volatility) may be a 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, because raising the equity premium through the consumption-volatility channel pushes the term premium above the data before the premium reaches its target (Appendix G quantifies this trade-off directly). An extra factor could break this link, allowing the two risk premia to be matched separately and improving the fit of the bond moments without sacrificing the equity results. The same extension would relax another limitation of the one-factor structure: with a single state variable, every yield is a function of the same state, so the level and the slope of the yield curve move together, whereas in the data they vary partly independently. A further possibility is to 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 pricing episodes. The indirect evidence in Appendix H 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, among others.

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 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 2004–2024 target window, suggesting that liquidity conditions do not materially affect the empirical targets. More directly, the positive term premium is not a product of TIPS illiquidity. A liquidity premium affects mainly the level of TIPS yields, which is why the targets are defined as 10–5 yr differences that net it out. Recomputing those differences from yields adjusted with the d’Amico–Kim–Wei liquidity-premium estimates (d’ Amico *et al.*, 2018) moves the targets modestly higher rather than lower, from 0.457 to 0.535 for the spread and from 0.708 to 0.858 for the excess return (Table 6).

The dividend-growth volatility σ_D warrants a brief note on how it is measured. Because D_t in (13) has independent increments, σ_D is the standard deviation of log dividend growth at an annual horizon, and it is measured directly at that horizon. It cannot be recovered by scaling up higher-frequency growth. The monthly dividend series of Shiller (2015) interpolates annual data, so monthly growth is close to deterministic and strongly serially correlated, and a $\sqrt{12}$ scaling of it understates annual volatility by roughly a factor of four. Both the mean and the standard deviation of real dividend growth are reported across sample periods in Table 4.

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). Dividend-growth volatility is the standard deviation of real dividend growth across calendar years, the empirical counterpart of σ_D .

Moment	Model	1990–2024	1999–2024	1926–2024
Mean P/D ratio	60.6	53.85	57.97	34.51
Median 2-yr std log(P/D)	0.060	0.067	0.074	0.096
Mean real div. growth (%)	—	2.89	3.41	1.97
Std. annual real div. growth (%)	—	7.16	8.20	10.40
Mean real stock return (%)	4.64	8.42	6.20	8.16
Std. log real stock return (%)	14.86	12.50	13.08	15.30
Mean equity premium (%)	3.02	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	2004–2024	2005–2024
Mean 10yr–5yr yield spread	0.422	0.457	0.440
Mean excess return 10yr/5yr	0.754	0.708	0.576

Table 6: TIPS targets under the d’Amico–Kim–Wei liquidity adjustment. Adjusted yields subtract the DKW liquidity-premium estimates (Fed Board updated series) from the fitted TIPS yields at the five- and ten-year maturities. Both moments are then recomputed exactly as in the baseline. ℓ_5 and ℓ_{10} are the mean liquidity premia over the window. All entries in % p.a.

Window	Spread 10–5		Excess return 10–5		$\ell_5 - \ell_{10}$
	raw	adjusted	raw	adjusted	
2004–2024 (target window)	0.457	0.535	0.708	0.858	0.078
2005–2024 (subsample)	0.440	0.513	0.576	0.720	0.073
1999–2024 (extended)	0.436	0.578	0.732	0.739	0.142

B Functional Forms for Consumption Dynamics

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

$$d \log C_t = \mu_c(x_t) dt + \sigma_c(x_t) dW_{ct}, \quad (31)$$

where x_t is a scalar latent state variable and W_{ct} is a standard Brownian motion.

The volatility schedule is a capped logistic function of the state,

$$\sigma_c(x) = \sigma_{\min} + (\sigma_{\max} - \sigma_{\min}) L\left(\frac{x - m}{w}\right), \quad L(u) = \frac{1}{1 + e^{-u}}, \quad (32)$$

which is strictly increasing and bounded on both sides. High- x states are high-volatility states, and the floor and cap keep the risk-adjusted dynamics well behaved at extreme state values (in particular, the cap bounds the state drift under the dividend measure; without the cap the P/D integral would diverge).

The real-rate schedule is a logistic with the same location m and width w as the volatility schedule (32),

$$r(x) = r_{\min} + (r_{\max} - r_{\min}) L\left(\frac{x - m}{w}\right), \quad (33)$$

with $(r_{\min}, r_{\max})$ calibrated. This makes $r(x)$ monotone by construction, flat near $r_{\min}$ in low-volatility states and rising through the same band of states over which volatility rises.³¹

The consumption drift is then implied by the CRRA short-rate identity rather than parameterized separately:

$$\mu_c(x) = \frac{r(x) - \rho}{\gamma} + \frac{\gamma}{2} \sigma_c(x)^2. \quad (34)$$

³¹Freeing the width of the rate transition relative to the volatility transition was tested and rejected by the data (Appendix D.3).

C Intermediary Consumption Volatility

This appendix documents two measurements used in Section 3: recovering the consumption volatility schedule from dealer equity volatility, and validating the sign and magnitude of ρ_{cx} against the dealer return–volatility correlation.

Persistence and the volatility schedule. This paragraph documents the measurement inputs used in Section 3. The data are the daily value-weighted equity returns of NY Fed primary dealer holding companies (He *et al.*, 2017), 2000–2018, summarized by their one-year rolling volatility (Panel A of Table 17).³² The monthly AR(1) coefficient of the rolling volatility series implies an annual persistence of $\phi = 0.890$. Given ϕ , the parameters $(\sigma_{\min}, m, w)$ of the logistic schedule (32) are fitted to the dealer-volatility quantiles through the wealth mapping, as described in Section 3. Because the wedge is state-dependent and the calibrated parameters depend in turn on the schedule, measurement and calibration form a joint fixed point: the wedge is computed at the current calibration, the schedule is refitted, the free parameters are recalibrated, and the cycle repeats. The iteration converges within two rounds, with the schedule parameters changing by less than 0.005 in the final round. The resulting fit is close at every targeted point of the σ_N distribution: 5th percentile 9.5% (data 9.6%), median 19.1% (data 19.5%), 95th percentile 55.5% (data 56.4%), mean 23.9% (data 23.1%). The consumption-volatility schedule this implies is reported in Section 3: 5th percentile 6.1%, median 10.5%, 95th percentile 33.4%, mean 13.8%.

The direction of this correction is worth noting. The wedge exceeds one everywhere, so the model-implied step of the measurement only ever removes volatility, and it removes proportionally the most around the middle of the stationary distri-

³²Specifically, the `intermediary_value_weighted_investment_return` series from the He *et al.* (2017) dataset is used. This is the value-weighted investment return to a portfolio of NY Fed primary dealers’ publicly-traded holding companies. Unlike the intermediary capital risk factor, this portfolio is tradable. Per the dataset’s documentation, He, Kelly, and Manela source the daily series entirely from Bloomberg.

bution, where the state spends most of its time. High consumption volatility is what generates the equity and bond premia, so the mapping withholds volatility in the states that matter most for them. A measurement that leaned on the model to flatter the fit would have to add volatility instead. Trusting the model’s wedge therefore biases the measurement against the volatility mechanism, not in its favor, and the unrestrained calibration of Appendix D.4 confirms that the asset-price moments would prefer the volatility the wedge removes.

The return–volatility correlation. This paragraph assesses whether the assumed perfect correlation of -1 between dealer returns and innovations to dealer volatility estimates, ρ_{Nx} , is reasonable. At first glance, this correlation looks weak in the data. Computed from one-month returns and volatility changes it is only -0.22 . This is an artifact of the horizon. The correlation grows steadily over longer windows and reaches -0.73 at twelve months (Table 7). Short-horizon estimates are biased toward zero, because the series is a one-year rolling estimate and is therefore an imperfect measure of volatility at short horizons. To correct for this bias, the same estimator is applied to wealth returns simulated from the model (the model counterpart of the dealer series, combining consumption shocks with the valuation term $d \log v(x_t)$) at known values of the underlying correlation. The model columns of Table 7 show the two features that matter. The bias is large. Even at the perfect-correlation limit the simulated estimator reaches only about half of the true correlation at the longest horizon. And the estimator barely varies with the underlying ρ_{cx} . The measured correlations also match or exceed the simulated ones at every horizon, even against the perfect-correlation limit itself. Two conclusions follow. First, the structural value $\rho_{cx} = -1$ does not overstate the measured correlations. The data ask for at least as much return–volatility comovement as the structure supplies. Second, the extra strength of the empirical correlations, beyond what the wedge already explains, points to sources of

comovement (such as jumps) that a correlated diffusion lacks, consistent with the discussion of the equity premium in Appendix G.

Table 7: Return–volatility correlations by horizon. “Data” is the correlation between NY Fed primary-dealer equity returns and changes in their one-year rolling volatility, computed over non-overlapping windows of the stated horizon (2000–2018, 217 monthly observations). The model columns apply the identical estimator to wealth returns simulated from the calibrated model at two values of the underlying consumption–state correlation. Means across 400 simulated samples, standard deviations in parentheses. The estimator is strongly attenuated at short horizons and nearly flat in ρ_{cx} , and the data lie at or beyond the perfect-correlation column at every horizon.

Horizon	Data	Model, $\rho_{cx} = -1$	Model, $\rho_{cx} = -0.3$
1 month	−0.22	−0.07 (0.08)	−0.04 (0.08)
3 months	−0.37	−0.18 (0.11)	−0.12 (0.13)
6 months	−0.48	−0.29 (0.15)	−0.19 (0.18)
12 months	−0.73	−0.52 (0.19)	−0.33 (0.25)

D Calibration Procedure and Robustness

This appendix documents the numerical targeting procedure, the leave-one-out robustness of the baseline calibration, and the test of the shared-transition restriction referenced in the main text.

D.1 Numerical procedure

After the directly set and the measured parameters are fixed (Section 3), the free vector is $\theta = (r_{\min}, r_{\max})$. Bounded parameters are mapped to the real line by a logit transformation, and the free vector is estimated by minimizing the inverse-variance weighted distance between eight simulated moments and their empirical targets,

$$\hat{\theta} = \arg \min_{\theta} \sum_{k=1}^8 w_k (m_k(\theta) - \hat{m}_k)^2, \quad w_k = 1/\hat{\sigma}_k^2, \quad (35)$$

where $m_k(\theta)$ is the model-implied value of moment k , $\hat{m}_k$ its empirical target, and $\hat{\sigma}_k$ the cross-simulation standard deviation of m_k , estimated from simulations (2,000 paths of 35 years). For each candidate θ the model is solved exactly (the bond pricing PDE and the dividend-strip system) and moments are computed from simulated paths using common random numbers, so that objective differences across parameter values are not driven by simulation noise. The optimization is run from a 3×3 grid of starting values for the two rate parameters, with each starting point optimized by three algorithms (Nelder–Mead, subplex, and CMA-ES, as implemented in the Julia packages Optim.jl (Mogensen and Riseth, 2018), NLOpt (Johnson, 2007), and CMAEvolutionStrategy.jl, respectively), for twenty-seven runs in all, summarized in Table 8. No restriction is placed on the sign of the consumption drift, and the two lower $r_{\min}$ values start the optimization with negative drift in calm states. All runs converge to the same interior optimum, agreeing on $r_{\min}$ and $r_{\max}$ to within rounding and on the objective to two decimals, so the solution is not sensitive to initialization or algorithm choice. At the optimum the drift is positive throughout the state space, so the rising rate schedule and the positive consumption drift are results of the fit rather than imposed features. The measured parameters are held fixed only within an iteration. Dealer volatility identifies $\sigma_c(x)$ only through a wedge that the model itself supplies, so the schedule is not an outside datum that can be held fixed. It has to be re-measured at each new calibration. The minimization described here is therefore the inner step of the joint fixed point of Appendix C.

D.2 Leave-one-out robustness

Table 9 re-estimates the free parameters eight times, dropping one targeted moment at a time (seven moments versus two parameters, so each run remains heavily overidentified). The calibration is stable: no single moment drives it. Across all

Table 8: Calibration of the two rate parameters. The upper panel lists the nine starting points of the 3×3 grid over $(r_{\min}, r_{\max})$, each optimized by three algorithms (Nelder–Mead, subplex, and CMA-ES). All twenty-seven runs converge to the same interior optimum, reported in the lower panel together with the minimized objective (the 200-path optimization objective under common random numbers). The rates are in % p.a. The correlation is fixed at $\rho_{cx} = -1$ (Section 3) and the time-preference rate at $\rho = 0.01$. The sign of the consumption drift is unrestricted, and starts with $r_{\min} \leq 0$ begin with negative drift in low-volatility states. The drift at the optimum is positive.

Start		
$r_{\min}$	$r_{\max}$	
−0.60	6.60	
−0.60	7.50	
−0.60	9.00	
+0.00	6.60	
+0.00	7.50	
+0.00	9.00	
+0.60	6.60	
+0.60	7.50	
+0.60	9.00	
Converged (all runs)		
$r_{\min}$	$r_{\max}$	Objective
+0.61	7.41	27.14

eight drops $r_{\min}$ stays within $[0.004, 0.009]$ and $r_{\max}$ within $[0.065, 0.082]$, against a baseline of $(0.006, 0.074)$. The equity premium is the instructive case. With the premium dropped, the parameters barely move and the model still delivers 2.9%, so the calibration concedes the premium rather than distorting itself to chase it. The documented miss is therefore a property of the model rather than of the weighting, and its sources are examined in Section 4.2 and Appendix G.

D.3 Testing the shared-transition restriction

In the baseline, the rate logistic $r(x)$ in (33) and the volatility logistic $\sigma_c(x)$ in (32) share the same center m and width w , so both transitions occur over the same band

Table 9: Leave-one-out robustness. Each row drops one targeted moment and re-optimizes the two free parameters against the remaining seven, with unchanged weights and seed. The correlation is fixed at $\rho_{cx} = -1$ throughout, as in the baseline, and the time-preference rate at $\rho = 0.01$. Baseline: $r_{\min} = 0.006$, $r_{\max} = 0.074$.

Dropped moment	$r_{\min}$	$r_{\max}$
Equity excess return	0.008	0.072
Mean P/D ratio	0.004	0.065
Mean real risk-free rate	0.009	0.079
Std. stock return	0.004	0.082
P/D volatility (2-yr)	0.007	0.073
TIPS 10–5 spread	0.006	0.072
TIPS 10–5 excess ret.	0.006	0.074
Return pred. slope	0.006	0.074

of states. To ask whether a different rate-transition width would be justified, a single parameter k is introduced to scale it, $r(x) = r_{\min} + (r_{\max} - r_{\min})L((x - m)/(kw))$, so the maximum rate slope scales as $1/k$. The relaxation is motivated by the model’s overshooting of yield-volatility statistics, and a flatter rate schedule might be expected to tame them. The opposite turns out to be true. Widening the transition drags rate variation out of the rarely visited tail into frequently visited states, so realized yield volatility rises even as the maximum slope falls. The free parameters are re-optimized at each $k \in \{1, 1.5, 2, 3, 4\}$ under the identical objective and seed, and the profile-minimizing k is adopted only if it improves the objective by more than the χ^2_1 5% critical value (3.84) and moves the Fama–Bliss standard-error statistics toward their data range. Both criteria fail. The profile minimum is $k = 1$ itself, and the objective rises monotonically, to 28.8 at $k = 1.5$, 34.8 at $k = 2$, 51.8 at $k = 3$, and 66.6 at $k = 4$, against 27.1 at $k = 1$. And the yield-volatility statistics move away from the data at every $k > 1$, as the mechanism above predicts, with the two-year Fama–Bliss standard error nearly tripling already at $k = 1.5$. The shared transition is therefore retained, and the yield-volatility overshoot is reported as a residual miss rather than repaired with an extra degree of freedom. Understanding the relationship between the rate and volatility schedules

more precisely remains an avenue for further research, particularly given that real short-term rates are also shaped by monetary policy.

D.4 Freeing the volatility schedule

It is interesting to see how well the model can do when the volatility schedule is not tied to the dealer data and the asset-price moments are allowed to determine it themselves. This subsection reports that exercise. The schedule’s floor, location, and width, $(\sigma_{\min}, m, w)$, the persistence φ_x , and the correlation ρ_{cx} , otherwise fixed at -1 by the wealth interpretation, are all freed and re-optimized jointly with the two rate parameters against the same eight moments, under the identical weights and seed. That leaves seven free parameters against eight moments. If the volatility schedule is unrestrained there is no reason to hold its persistence fixed, so φ_x is freed along with it. Because the free variant fits σ_c to asset prices directly, with no wealth inversion, its cap is imposed in consumption-volatility space at 0.75, just above the dealer-data maximum (asset prices do not identify the cap, and a free cap drifts to values at which the dividend strips fail to converge). Table 10 reports the estimated parameters alongside the baseline.

Three results follow.

First, the perfect-correlation restriction is corroborated, not merely tolerated. With ρ_{cx} entirely free, the fit selects -1 from every starting point. Given the choice, the asset-price moments place the correlation exactly where the wealth foundation of Section 2.7 puts it, so relaxing the restriction buys nothing.

Second, the central finding is intact, and the baseline turns out to be the conservative reading. With the floor, location, width, and persistence all determined by the asset-price moments, the fitted schedule implies a mean consumption volatility of 23.2% under the stationary distribution of the state, substantially larger than the 13.8% of the measured schedule. At the same time, both levels are

Table 10: Estimated parameters when the volatility schedule is freed. Each column re-optimizes the stated free parameters against the same eight moments, weights, and seed. In the baseline the persistence φ_x and the schedule parameters $(\sigma_{\min}, m, w)$ are measured from the dealer data through the wealth mapping and the correlation is fixed at -1 . The free column lets the asset-price moments determine all seven. The cap conventions differ. The baseline cap $\sigma_{\max} = 0.623$ is solved from $\max_x \sigma_N = 0.75$ in wealth-volatility space, while the free variant fits σ_c directly and its cap is imposed at 0.75 in consumption-volatility space, as asset prices do not identify it.

Parameter	Symbol	Baseline	Free
Volatility schedule and persistence			
State-variable persistence	φ_x	0.890	0.943
Consumption vol. (floor)	$\sigma_{\min}$	0.056	0.142
Volatility transition (center)	m	0.691	1.500
Volatility transition (width)	w	0.293	0.771
Consumption vol. (cap)	$\sigma_{\max}$	0.623	0.750
Rate schedule and preferences			
Consumption–state correlation	ρ_{cx}	-1	-1
Real rate (low volatility)	$r_{\min}$	0.006	-0.016
Real rate (high volatility)	$r_{\max}$	0.074	0.095

an order of magnitude above the volatility of aggregate consumption. Table 11 reports the stationary quantiles. The free and measured schedules agree closely in the right tail, placing the ninety-fifth percentile at 34.4% against 33.4%, and diverge in calm states, where the free floor is 14% and the free median 22%, against 5.6% and 10.5% measured. Left to themselves, in other words, the asset-price moments would earn the equity premium by keeping consumption volatility high even in tranquil times. The dealer data say tranquil times are genuinely tranquil, and the measurement forces the model to live with the premium that remains. Asset prices alone would demand more volatility than the measurement supplies, so the measured baseline understates rather than overstates what the mechanism could deliver. As discussed in Appendix C, the wedge exceeds one everywhere, so trusting it biases the measurement against the volatility mechanism that generates the premia, not in its favor.

Freeing the schedule, its persistence, and the correlation lowers the objective from 26.03 to 12.88.³³ Some decline is to be expected, since a minimized objective can only fall when parameters are freed. About nine of the thirteen points come from the equity premium, which rises from 3.0% to 4.6% against a target of 8.4%. The remainder comes from the risk-free rate and the stock return volatility, partly offset by a worse fit of the TIPS and price-dividend moments. The exercise is not a horse race between two calibrations of equal standing. The measured schedule is pinned by the dealer-volatility series, information the asset-price objective never sees, and the comparison quantifies what that measurement discipline costs in fit.

Third, the improvement in the eight targeted moments is bought with counterfactual behavior in the moments that were not targeted. The Fama-Bliss slopes rise to 4.3–4.9 at maturities of two to five years, against data values of 1.3–1.8 and baseline values of 2.6–2.7; the Campbell-Shiller slopes fall to -7.2 and below, against data values between -1.0 and -2.8 and baseline values near -4.2 ; and the Newey-West standard errors, already an acknowledged overshoot at the baseline, rise further. The extra calm-state volatility that raises the premium also drives term-structure dynamics far from the data. The free fit also shows boundary strain the baseline does not have, since it pins the transition location at the edge of its admissible range.

Even with the volatility process entirely free, the model delivers barely more than half the empirical premium while respecting the remaining moments. The premium gap documented in Appendix G is therefore not an artifact of the measurement discipline. It survives the most permissive calibration the model admits.

³³Objectives in this comparison are evaluated with 2,000 simulated paths, so the baseline value is 26.03 rather than the 27.1 reported in Section 3, which uses the 200 paths of the optimization.

D.5 Raising the risk-free-rate target

As noted in Section 3, the model’s $r(x)$ is the marginal investor’s own shadow real rate, while the empirical target is the one-month Treasury bill rate less inflation. Treasury securities carry a convenience yield that a leveraged marginal investor does not earn, estimated at 40–70 basis points by Krishnamurthy and Vissing-Jorgensen (2012), and an investor constrained by its own balance sheet faces on top of that the shadow cost of its funding capacity (Gârleanu and Pedersen, 2011). Both wedges place the investor’s shadow rate above the observed bill rate. Because this target enters the objective through its first moment alone, such a wedge is exactly a shift in a single number, and the calibration can simply be repeated with the target raised accordingly.

Table 12 reports the calibration re-run with the target raised by 40, 70, and 100 basis points. The two free parameters $r_{\min}$ and $r_{\max}$ move only slightly, and the remaining seven moments are almost untouched. The model’s own mean real rate rises by less than the shift, so the gap against the target narrows as the wedge is granted, while the fit of every other moment is essentially preserved.³⁴ The calibration is therefore not an artifact of identifying the bill rate with $r(x)$.

D.6 The sign of the consumption–state correlation

The wealth foundation fixes the magnitude $|\rho_{cx}| = 1$, since a single innovation drives both wealth and the state, but the negative sign is a modeling choice. To check that it is not merely a normalization, the entire joint measurement–calibration fixed point is re-run with $\rho_{cx} = +1$. The volatility schedule is re-measured under the positive-sign wealth mapping and the two rate parameters are re-optimized to the same eight moments. The baseline restricts the rate schedule to rise with

³⁴The minimized objective also falls across these runs, but that is not evidence of a better fit, since each run measures distance to a different target and the shift moves that target toward what the model already delivers.

the state ($r_{\max} > r_{\min}$). Under $\rho_{cx} = +1$ that restriction makes the present value of the dividend claim diverge, because a positive correlation turns equities into a hedge and lowers their discount rate, so the restriction is relaxed here and the schedule is free to decrease.

Table 13 reports the outcome. The positive sign is admissible, but only with an inverted rate schedule, the real rate falling from +5.8% in calm states to -5.0% in the most volatile states. Even then the fit is about three times worse than the baseline, a weighted objective of 72.3 against 26.0. The excess volatility of stock prices collapses, with the two-year volatility of $\log(P/D)$ falling to 0.016 from 0.060 and the standard deviation of returns to 8.0% from 14.9%, and the TIPS excess return drops to 0.30% from 0.75%. The equity premium is only somewhat lower, 2.3% against 3.0%. The negative sign is therefore preferred on grounds of both fit and economic plausibility, since the positive sign requires a real rate that turns sharply negative exactly when volatility is highest and gives up the excess-volatility and term-premium behavior that motivates the model.

E Simulation

Figure 3 presents a simulated forty-year path of the baseline model's state variable and the asset prices implied by it.




Figure 3: Variables based on a simulated path for forty years under the baseline parameterization. The panels show the consumption drift μ_c , consumption volatility σ_c , the risk-free rate r , the price–dividend ratio, the expected stock excess return, and the 10-year real bond yield together with its term premium.

Table 11: Freeing the volatility schedule: implied moments. Each column reports the moments implied by the corresponding calibration in Table 10, under the same weights and seed. The first eight rows are the targeted moments. The moments below them are untargeted, with Fama–Bliss slopes reported by maturity and Campbell–Shiller slopes by horizon. Volatility quantiles are stationary quantiles of $\sigma_c(x)$ and summarize the schedule rather than an asset-price moment. The baseline column reports the measured (wealth-inverted) schedule. No data entry is given for the volatility quantiles because the dealer-volatility data live in wealth-volatility space, which the baseline matches by construction through the mapping of Section 3. The Data column gives the empirical targets. Stock moments are over 1990–2024 (Shiller, 2015) and TIPS moments over 2004–2024 (Gürkaynak *et al.*, 2010). Fit objectives and their statistical assessment are reported in the text. The free column and the first column of Table 16 free the same seven parameters over the same box and differ only in the target set, all eight moments here against the equity premium alone there.

Moment	Data	Baseline	Free
Targeted moments			
Equity excess return (%)	8.38	3.02	4.62
Mean P/D	53.9	60.6	62.4
Median 2-yr std $\log(P/D)$	0.067	0.060	0.056
Std. stock return (%)	12.50	14.86	13.45
Mean real risk-free rate (%)	0.03	1.62	0.00
TIPS 10–5 yr spread (%)	0.46	0.42	0.52
TIPS 10–5 yr excess return (%)	0.71	0.75	1.06
Return pred. slope $\hat{\beta}_1$	0.21	0.25	0.20
Untargeted moments			
10-yr bond excess return (%)	2.55	1.81	2.09
Fama–Bliss slopes ($\beta = 1$ under EH)			
2-year	1.33	2.59	4.32
3-year	1.71	2.64	4.48
4-year	1.84	2.65	4.66
5-year	1.69	2.62	4.87
Campbell–Shiller slopes ($\beta < 0$ violates EH)			
3-month	-1.03	-4.03	-7.23
6-month	-1.16	-4.06	-7.29
1-year	-1.41	-4.13	-7.41
2-year	-1.92	-4.26	-7.66
4-year	-2.83	-4.48	-8.19
Volatility schedule $\sigma_c(x)$, stationary quantiles (%)			
5th percentile	—	6.1	16.6
Median	—	10.5	21.8
95th percentile	—	33.4	34.4
Mean	—	13.8	23.2

Table 12: Sensitivity of the calibration to the level of the mean real-rate target. Each column re-runs the joint calibration with the mean real-rate target raised by the stated amount, holding the weights, the other seven moments, the measured schedule, and the seed fixed. Only the first moment of the rate target moves. Rates and returns in %.

Quantity	+0 bp	+40 bp	+70 bp	+100 bp
Mean real-rate target (%)	0.03	0.43	0.73	1.03
$r_{\min}$ (%)	0.61	0.66	0.71	0.75
$r_{\max}$ (%)	7.41	7.48	7.53	7.60
Model mean real rate (%)	1.62	1.68	1.72	1.77
Equity premium (%)	3.02	3.00	2.99	2.98
Mean P/D	60.6	59.2	58.1	56.9
Std. stock return (%)	14.86	14.78	14.71	14.65
Median 2-yr std $\log(P/D)$	0.060	0.060	0.060	0.059
TIPS 10–5 yr spread (%)	0.42	0.42	0.42	0.42
TIPS 10–5 yr excess return (%)	0.75	0.76	0.76	0.76
Return pred. slope $\hat{\beta}_1$	0.250	0.250	0.250	0.251

Table 13: The sign of ρ_{cx} . The $\rho_{cx} = +1$ column re-runs the joint measurement–calibration fixed point with the sign flipped, the volatility schedule re-measured under the positive-sign wealth mapping, and the two rate parameters re-optimized to the same eight moments with the increasing-rate restriction relaxed. Objectives are evaluated with 2,000 simulated paths, so the baseline value is 26.0 rather than the 27.1 of Section 3, which uses the 200 paths of the optimization. Rates and returns in %.

Moment	Data	$\rho_{cx} = -1$	$\rho_{cx} = +1$
Equity excess return (%)	8.38	3.02	2.26
Mean P/D	53.9	60.6	48.9
Median 2-yr std $\log(P/D)$	0.067	0.060	0.016
Std. stock return (%)	12.50	14.86	7.99
Mean real risk-free rate (%)	0.03	1.62	2.77
TIPS 10–5 yr spread (%)	0.46	0.42	0.42
TIPS 10–5 yr excess return (%)	0.71	0.75	0.30
Return pred. slope $\hat{\beta}_1$	0.21	0.25	0.25
Weighted objective	—	26.0	72.3
$r_{\min}$ (%)	—	0.61	5.79
$r_{\max}$ (%)	—	7.41	–5.00
Rate schedule	—	increasing	decreasing

F Expectation-Hypothesis Regression Results

Tables 14 and 15 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 parameterization the sign of the estimates matches the data. FB betas are positive (and above unity) and CS betas are negative, both indicating a violation of the expectations hypothesis. The estimates are imprecise. Their cross-simulation dispersion is wide, the empirical slopes lie within it at every horizon, and the pattern varies considerably across parameterizations even though the violation itself is robust. As discussed in Section 4.3, these regressions are therefore a weak guide to the term-premium pattern. That pattern is read instead from the model-implied functions in Figure 2 of the main text, in particular the decomposition of the 10-year real yield into its risk-neutral and term-premium components.

Table 14: 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 35-year samples. Standard deviations across simulations in parentheses. $\beta = 1$ under the expectations hypothesis.

Maturity	$\bar{\beta}$	(SD)	Target
2-year	2.591	(2.151)	1.330
3-year	2.638	(2.493)	1.710
4-year	2.653	(2.897)	1.840
5-year	2.620	(3.392)	1.690

G Anatomy of the Equity Premium

The baseline delivers an equity premium of 3.0% against the 8.4% target, the one targeted moment on which the joint calibration falls clearly short. This appendix decomposes the shortfall and closes off cash-flow risk, the route that might seem to repair it.

Table 15: 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 35-year samples. Standard deviations across simulations in parentheses. $\beta < 0$ indicates violation of the expectations hypothesis.

Horizon	$\bar{\beta}$	(SD)	Target
3-month	-4.028	(3.300)	-1.030
6-month	-4.062	(3.394)	-1.160
1-year	-4.129	(3.585)	-1.410
2-year	-4.258	(3.981)	-1.920
4-year	-4.483	(4.831)	-2.830

G.1 Only targeting the premium

Table 16 traces the equity premium as the elements of the baseline calibration, the measured volatility schedule and the remaining target moments, are imposed one at a time. The first column re-calibrates the model with the equity premium as the sole target and the volatility schedule left entirely free. It then matches the premium, 8.6% against a target of 8.4%, so the premium is not beyond the model’s reach. But the fitted schedule puts consumption volatility at 20% even in low-volatility states, the TIPS 10–5 spread is 2.9% against 0.5% in the data, the mean real rate is +7.7% against a near-zero target, and the mean price–dividend ratio is 9 against 54.³⁵ This column also bounds what a higher risk aversion could buy. Because γ enters pricing only through the product $\gamma\rho_{cx}\sigma_c(x)$, raising it is equivalent, for pricing, to scaling the volatility schedule, and reaching the premium this way takes a price of risk roughly twice the measured one, at the cost of the real rate and the valuation moments just described. Risk aversion is therefore a margin the article declines to use rather than one the model lacks, and $\gamma = 2$ is retained. The second column keeps the premium as the sole target but imposes the measured schedule. The measurement alone cuts the deliverable premium to

³⁵The daggered entries in Table 16 are not pinned down: this column matches one moment with seven free parameters, so any of a wide range of schedules reaches the premium. What survives across all of them is the cost, which is what the column is for.

4.1%, and even that requires the rate ceiling at its admissible bound ($r_{\max} = 30\%$) and a price–dividend ratio of 20. In these columns the correlation ρ_{cx} is estimated rather than fixed, and the optimizer selects -1 from every starting point, which is the value the baseline imposes on structural grounds. The third column targets the six stock-market and risk-free-rate moments with no term-structure targets, and the premium settles at 3.0%; the fourth adds the two TIPS moments and barely moves any entry, the largest change being the price–dividend ratio, from 61.9 to 60.6. The decomposition locates the binding constraints precisely: the measured level of consumption volatility in calm states does the largest share of the capping, the equity valuation moments do the rest, and the real term structure — the natural suspect, since the equity premium, the price–dividend ratio, and the real term premium are all priced by the same $\gamma\rho_{cx}$ — turns out to bind hardly at all once the valuation moments are present. The shortfall is the visible cost of asking one factor to do all of these jobs honestly, as Appendix D.4 documents from the opposite direction.

G.2 Cash-flow risk is shut down by measurement

The baseline sets $\rho_{cD} = 0$ and $\sigma_D = 7.2\%$, the measured volatility of annual real dividend growth. Cash-flow risk cannot raise the premium per unit of return volatility. A dividend shock earns $\gamma\sigma_c\rho_{cD}$ per unit of volatility and a revaluation shock earns $\gamma\sigma_c|\rho_{cx}|$, and since $|\rho_{cx}| = 1$ at the baseline, dividend risk can at best tie it and never beat it. The only question is how much it dilutes, and the measured correlation settles that. Between dividend growth and dealer returns it runs between -0.07 and -0.19 across horizons, the opposite sign. Imposing $\rho_{cD} = -0.15$ lowers the premium to 2.7%, since cash-flow volatility that carries no premium crowds out revaluation volatility that was earning one. Even a counterfactually positive $+0.71$ reaches only 4.7%, and 5.2% with $\sigma_D = 0.10$, with mean price–dividend ratios of

Table 16: The equity premium as the elements of the baseline calibration are imposed one at a time. Each column re-optimizes the free parameters against the stated target set, under the baseline weights and seed. The first column additionally frees the volatility schedule (capped logistic form, no measurement discipline). The remaining columns impose the measured schedule of Section 3. The equity block comprises the six stock-market and risk-free-rate moments (excess return, mean and volatility of P/D, return volatility, mean real rate, predictability slope). The baseline adds the two TIPS moments. Table 11 completes the design: it frees the same seven parameters as the first column here, but targets all eight moments rather than the equity premium alone.

	Premium only, free σ_c	Premium only, measured σ_c	Equity block, no term struct.	Baseline	Data
Equity premium (%)	8.63	4.12	2.98	3.02	8.38
TIPS 10–5 (%)	2.92 [†]	1.16	0.41	0.42	0.46
Mean r_f (%)	+7.74 [†]	+4.93	+1.62	+1.62	+0.03
Mean P/D	8.8 [†]	20.1	61.9	60.6	53.9
ρ_{cx}	−0.89 [†]	−1.00	−1.00	−1 (fixed)	—

[†] Not determined by this column’s target set. The first column matches one moment with seven free parameters, so these entries are one point on a flat set rather than an estimate: across ten optimizations differing only in the optimizer’s random stream, every one reaches the premium (8.5–8.6%, objective zero to four decimals) while the mean real rate ranges from −2.5% to +12.0%, the price–dividend ratio from 6 to 41, and consumption volatility in calm states from 9% to 30%. The last column is determined: all twenty-seven runs of Table 8 reach the same optimum.

35 and 27 respectively against 54 in the data. The asset-price objective alone is nearly indifferent to that configuration (27.2 against the baseline’s 27.1), so it is the measurement of ρ_{cD} , not the fit, that rules it out.

Taken together, the re-calibrations and the closed cash-flow route locate the missing premium. Within one-factor diffusion dynamics and the measured volatility of the marginal investor, no cash-flow parameter reaches the target, and reaching it through the price of risk alone requires roughly twice the volatility the measurement supplies. The gap is evidence for amplification channels that the diffusion model deliberately omits, such as jumps and volatility feedback. This reading is consistent with the attenuation gap in the return–volatility correlations of Appendix C. Exploring these channels is left to future work.

H Asset-Price Evidence from the Recovered State

The state variable is measured from the volatility of dealer equity returns, so the prices it implies can be compared with the data as an out-of-sample check of the intermediary interpretation. The dealer return series is itself the price of a traded equity portfolio, so the measurement is not free of price information altogether. Only the variance of those returns enters, however, and not the level of any price, valuation ratio, or yield, and in particular none of the series it is compared with below. At each date the implied state x_t is recovered by inverting the model-implied wealth volatility $\sigma_N(x)$ at the observed dealer volatility, and bond and stock prices are then evaluated at that state. Figure 4 shows the dealer volatility series, the state recovered from it, and the implied price–dividend ratio, 10-year TIPS yield, and 10–5 yr TIPS spread against their empirical counterparts. Table 17 reports the corresponding correlations.

The recovered state comoves positively with equity valuations and the real term structure. The implied price–dividend ratio has a levels correlation of 0.32 with the observed ratio and the implied TIPS spread one of 0.47, discussed in Section 4.5. These are substantial for a state built without reference to either series. The level of the 10-year TIPS yield is the one series the mapping does not track, at 0.03 in levels and 0.02 in changes. The sample features a secular decline in real yield levels that a stationary one-factor state cannot generate, which is also why the article’s empirical targets are a maturity spread and an excess return rather than yield levels.

The match is of course far from exact, as one would expect from a one-factor model driven by a single measured series. Several broad movements are nonetheless reproduced. The model-implied price–dividend ratio falls sharply into the 2008 trough and recovers alongside its empirical counterpart, though its swings are considerably larger than those in the data. The implied TIPS spread stays close

to the average level of the observed spread but is much smoother, which is the one-factor structure showing through.

The inversion is well defined throughout the sample. Dealer volatility peaks at 72.6% near the 2008 financial crisis, below the maximum of σ_N at 75%, so the state is recovered exactly at every date, with the crisis mapping to a large positive realization of x_t .

Table 17: Dealer volatility and model-implied asset prices. 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. 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 the model-implied wealth volatility $\sigma_N(x)$ at each date using the empirical intermediary equity return volatility, then evaluating bond and stock prices at the implied state.

Panel A: Intermediary equity return vol (% p.a.)				
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	0.315		0.396	
10-yr TIPS yield	0.026		0.023	
10–5yr TIPS spread	0.471		0.152	

^aAR(1) on monthly series. Reported value is ϕ^{12} , the implied annual persistence.




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 the model-implied wealth volatility $\sigma_N(x)$ at each date using the empirical equity return volatility. The sample maximum (72.6%) lies below the hump maximum of $\sigma_N(x)$ at 75%, so the lower-branch inversion is exact at every date. 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). Shaded bands mark NBER recessions.

References

- ABRAHAM, 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.
- ANGELETOS, G.-M., COLLARD, F. and DELLAS, H. (2018). Quantifying confidence. *Econometrica*, **86** (5), 1689–1726.
- 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.

- BAUER, M. D. and HAMILTON, J. D. (2018). Robust bond risk premia. *The Review of Financial Studies*, **31** (2), 399–448.
- BEKAERT, G., ENGSTROM, E. and GRENADIER, S. R. (2010). Stock and bond returns with moody investors. *Journal of Empirical Finance*, **17** (5), 867–894.
- , HODRICK, R. J. and MARSHALL, D. A. (1997). On biases in tests of the expectations hypothesis of the term structure of interest rates. *Journal of Financial Economics*, **44** (3), 309–348.
- 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.
- , PFLUEGER, C. and VICEIRA, L. M. (2020). Macroeconomic drivers of bond and equity risks. *Journal of Political Economy*, **128** (8), 3148–3185.
- 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.
- COLLIN-DUFRESNE, P., JOHANNES, M. and LOCHSTOER, L. A. (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, **106** (3), 664–698.
- 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.
- DEN HAAN, W. J. (1995). The term structure of interest rates in real and monetary economies. *Journal of Economic Dynamics and Control*, **19** (5-7), 909–940.
- DONALDSON, J. B., JOHNSEN, T. and MEHRA, R. (1990). On the term structure of interest rates. *Journal of Economic Dynamics and Control*, **14** (3-4), 571–596.
- DUFFEE, G. R. (2002). Term premia and interest rate forecasts in affine models. *The Journal of Finance*, **57** (1), 405–443.
- (2013). Bond pricing and the macroeconomy. *Handbook of the Economics of Finance*, **2**, 907–967.
- 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.
- GÂRLEANU, N. and PEDERSEN, L. H. (2011). Margin-based asset pricing and deviations from the law of one price. *The Review of Financial Studies*, **24** (6), 1980–2022.
- GOMEZ-CRAM, R. and YARON, A. (2021). How important are inflation expectations for the nominal yield curve? *The Review of Financial Studies*, **34** (2), 985–1045.
- 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.
- HSU, A., LI, E. X. and PALOMINO, F. (2021). Real and nominal equilibrium yield curves. *Management Science*, **67** (2), 1138–1158.
- JAPPELLI, R., PELIZZON, L. and SUBRAHMANYAM, M. G. (2024). *Monetary Policy, the Yield Curve, and the Repo Market*. Tech. rep., SAFE Working Paper.

- JOHNSON, S. G. (2007). The NLOpt nonlinear-optimization package. <https://github.com/stevengj/nlopt>.
- JOYCE, M. A. S., KAMINSKA, I. and LILDHOLDT, P. (2012). Understanding the real rate conundrum: An application of no-arbitrage models to the UK real yield curve. *Review of Finance*, **16** (3), 837–866.
- KARGAR, M. (2021). Heterogeneous intermediary asset pricing. *Journal of Financial Economics*, **141** (2), 505–532.
- KATAGIRI, M. (2022). Equilibrium yield curve, the phillips curve, and monetary policy. *Journal of Money, Credit and Banking*, **54** (8), 2235–2272.
- KEKRE, R., LENEL, M. and MAINARDI, F. (2024). *Monetary Policy, Segmentation, and the Term Structure*. Tech. rep., National Bureau of Economic Research.
- KLIEM, M. and MEYER-GOHDE, A. (2022). (un) expected monetary policy shocks and term premia. *Journal of Applied Econometrics*, **37** (3), 477–499.
- KRAFT, H., SEIFERLING, T. and SEIFRIED, F. T. (2017). Optimal consumption and investment with epstein–zin recursive utility. *Finance and Stochastics*, **21**, 187–226.
- KRISHNAMURTHY, A. and VISSING-JORGENSEN, A. (2012). The aggregate demand for treasury debt. *Journal of Political Economy*, **120** (2), 233–267.
- 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.

- MERTON, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, pp. 247–257.
- 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).
- NAGEL, S. (2016). The liquidity premium of near-money assets. *The Quarterly Journal of Economics*, **131** (4), 1927–1971.
- 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.
- RUDEBUSCH, G. D. and SWANSON, E. T. (2008). Examining the bond premium puzzle with a DSGE model. *Journal of Monetary Economics*, **55**, S111–S126.
- and — (2012). The bond premium in a DSGE model with long-run real and nominal risks. *American Economic Journal: Macroeconomics*, **4** (1), 105–43.
- 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? *American Economic Review*, **71**, 421–436.

- (2015). *Irrational Exuberance: Revised and Expanded Third Edition*. Princeton, NJ: Princeton University Press.
- TSAI, J. (2015). Rare disasters and the term structure of interest rates. *Available at SSRN 2669992*.
- 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.
- (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, **68** (3), 987–1035.
- 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.

Internet Appendix

“High Volatility is the Key: A Joint Treatment of Asset Pricing Puzzles”

This Internet Appendix consists of three parts. Section IA.I provides a review of the literature on the bond premium puzzle. Section IA.J presents a self-contained mathematical framework for the model variations analyzed in this appendix. Section IA.K documents yield and term-premium results for all 37 model variations, together with a model index table. This appendix is adapted from an earlier working-paper version of this article and is included here because the main text makes reference to it. Its model variations and calibrations predate the baseline calibration of the main text and share no parameter values with it. The appendix documents qualitative conclusions, in particular that standard models do not generate positive real term premia across a wide range of specifications, and its numbers should not be compared with those of the main text.

IA.I Literature on the Bond Premium Puzzle

This section reviews available explanations of the bond premium puzzle.³⁶ One of the first papers to address this was Backus *et al.* 1989. Utilizing a consumption-based asset-pricing model of an endowment economy, they discovered the model’s inability to yield significant positive term premia. Subsequent studies by Donaldson, Johnsen and Mehra 1990 and Den Haan 1995 further indicated that standard real business cycle models also could not resolve the puzzle. Rudebusch and Swanson 2008 incorporated an external habit into DSGE models but found that the bond premium puzzle remains. Duffee 2013 showed that basic properties of nominal

³⁶Rudebusch and Swanson 2008 provides a good summary of this extensive literature.

yields cannot be explained macroeconomically. Also, Duffee 2002 shed light on the challenges of fitting both interest rate and term premium dynamics within affine models.

Next, a series of papers provided explanations focused on nominal term premia. Notably, Piazzesi and Schneider 2006 showed that parameter uncertainty in a model where inflation brings bad news about future consumption growth can produce positive nominal term premia.³⁷ Gabaix 2012 and Tsai 2015, following Rietz 1988 and Barro 2006, showed that positive nominal term premia can be explained if inflation is on average high during consumption disasters. Bansal and Shaliastovich 2013, following Bansal and Yaron 2004, demonstrated that the risk premium of a nominal bond can be positive when inflation is correlated with consumption trend. Rudebusch and Swanson 2012 used a similar model within a DSGE framework that generates positive nominal term premia, but real term premia are negative. Gomez-Cram and Yaron 2021 also used a long-run risk model focused on nominal term premia.

Alternatively, some articles consider real term premia. Katagiri 2022 explored a model with monetary policy in which consumption changes can be negatively correlated with consumption trends and risk aversion is very high. Ellison and Tischbirek 2021 used a beauty-contest mechanism as in Angeletos, Collard and Dellas 2018, in which agents anticipate others' expectations, generating positive term premia.

Using an approach similar to the main article, some papers deviate from the representative agent model. Vayanos and Vila 2021 suggested that term premia are generated by arbitrageurs interacting with preferred-habitat investors. Kekre *et al.* 2024 built on Vayanos and Vila 2021 and showed that arbitrageur portfolio characteristics have important implications for the sign of term premia. Jappelli, Pelizzon

³⁷Collin-Dufresne, Johannes and Lochstoer 2016 introduced a model with Bayesian learning of parameters that generates negative term premia.

and Subrahmanyam 2024 extended this by integrating the repo market. Schneider 2022 showed that positive term premia can arise in models with heterogeneous agents exhibiting different risk attitudes and intertemporal substitution preferences.³⁸ Finally, returning to representative-agent models, Wachter 2006 showed that term premia can be positive and time-varying in a model with external habit following Campbell and Cochrane 1999. Kliem and Meyer-Gohde 2022 and Hsu, Li and Palomino 2021 used the same mechanism within DSGE models. Campbell, Pflueger and Viceira 2020 also used a habit model to explain the time-variability of term premia. More generally, Lettau and Wachter 2011 and Bekaert, Engstrom and Grenadier 2010 showed that positive and time-varying term premia can be obtained within the class of models with time-varying effective risk aversion. To the authors' knowledge, time-varying risk aversion is the only mechanism in the literature that achieves positive and time-varying term premia within a rational representative agent model.

IA.J Mathematical Framework

A consumption-based framework in continuous time accommodates a range of model variations. The framework is built upon three components: (1) an exogenous consumption process, (2) a utility specification, and (3) a process for the state variable. The state variable determines the state of the economy, and is connected to either the consumption drift (CD), the consumption volatility (CV), or the external habit. Utility is either time-separable (TSU) or recursive (RU), the latter following Duffie and Epstein 1992.

³⁸Schneider 2022 also uses effective time-varying risk aversion.

IA.J.1 Consumption process

Consumption (C_t) follows the stochastic process:

$$d \log(C_t) = dc_t = \mu_{ct} dt + \sigma_{ct} dW_{ct} \quad (\text{IA.36})$$

where μ_{ct} denotes the consumption drift at time t and σ_{ct} is the consumption volatility, multiplying the standard Wiener process dW_{ct} .

IA.J.2 Utility

Lifetime utility takes the following form depending on the specification:

$$\underbrace{U_0 = E_0 \int_0^\infty e^{-\rho t} u(C_t, S_t) dt}_{\text{TSU}}, \quad \underbrace{V_0 = E_0 \int_0^\infty F(C_t, V_t) dt}_{\text{RU}} \quad (\text{IA.37})$$

where ρ is the time preference parameter. In the TSU case, flow utility u depends on consumption and potentially on the surplus consumption ratio S_t (relevant only in the habit model; set to 1 otherwise). In the RU case, F is the Duffie–Epstein aggregator depending on consumption and the current value function V_t . The functional forms are:

$$u(C_t, S_t) = \frac{(C_t S_t)^{1-\gamma} - 1}{1-\gamma} \quad (\text{TSU})$$

$$F(C_t, V_t) = \frac{\rho(1-\gamma)V_t}{1-1/\psi} \left(\left(\frac{C_t}{((1-\gamma)V_t)^{-1/(1-\gamma)}} \right)^{1-1/\psi} - 1 \right) \quad (\text{RU}) \quad (\text{IA.38})$$

Here γ is risk aversion (equal to the inverse of the IES in TSU), and ψ is the IES in the RU case.

IA.J.3 State variable

The state variable x_t follows a mean-reverting process:

$$dx_t = -\log(\phi)(\mu_{x0} - x_t) dt + \sigma_{xt} dW_{xt} \quad (\text{IA.39})$$

where μ_{x0} is the steady state and $\phi \in (0, 1)$ governs mean reversion. The Wiener process dW_{xt} has correlation ρ_{cx} with dW_{ct} . Table IA.18 shows how each model variation maps the state variable into the consumption and state processes.

Table IA.18: **State variable dependence for each model variation**

Model variation			
TSU-CD:	$\mu_{ct} = \mu_{c0} + x_t$	$\sigma_{xt} = \sigma_{x0}$	$\mu_{x0} = 0$
TSU-CV:	$\sigma_{ct} = \sigma_{c1}x_t$	$\sigma_{xt} = \sigma_{x1}\sqrt{x_t}$	$\mu_{x0} = 1$
TSU-Habit:*	$S_t = \bar{S}e^{x_t}$	$\sigma_{xt} = \sigma_{ct}\lambda(x_t)$	$\mu_{x0} = 0$
RU-CD:	$\mu_{ct} = \mu_{c0} + x_t$	$\sigma_{xt} = \sigma_{x0}$	$\mu_{x0} = 0$
RU-CV:	$\sigma_{ct} = \sigma_{c1}\sqrt{x_t}$	$\sigma_{xt} = \sigma_{x1}\sqrt{x_t}$	$\mu_{x0} = 1$
TSU-HCV:	$\sigma_{ct} = \sigma_{c1}x_t$	$\sigma_{xt} = \sigma_{x1}\sqrt{x_t}$	$\mu_{x0} = 1$
Arb-IN:	$\mu_{ct} = \mu_{c1}x_t^{1/4}, \sigma_{ct} = \sigma_{c1}\sqrt{x_t}$	$\sigma_{xt} = \sigma_{x1}\sqrt{x_t}$	$\mu_{x0} = 1$
Arb-DP:	$\mu_{ct} = \mu_{c1}x_t^{3/2}, \sigma_{ct} = \sigma_{c1}x_t^{2/3}$	$\sigma_{xt} = \sigma_{x1}\sqrt{x_t}$	$\mu_{x0} = 1$

* Following Campbell and Cochrane 1999 and Wachter 2006:

$$\sigma_{xt} = \sigma_{ct}\lambda(x_t) = \begin{cases} \sigma_{ct}\left(\frac{\sqrt{1-2x_t}}{\bar{S}} - 1\right) & \text{if } x_t < \frac{1-\bar{S}^2}{2} \\ 0 & \text{if } x_t \geq \frac{1-\bar{S}^2}{2} \end{cases}, \quad \bar{S} = \sqrt{\frac{\gamma}{-\log(\phi) - b/\gamma}}$$

In variations where consumption volatility σ_{ct} is proportional to the state variable, the steady state is set at $x_t = 1$ (with $x_t > 0$ almost surely), ensuring positivity of σ_{ct} .

IA.J.4 Stochastic discount factor

Time-separable utility. The SDF is the derivative of flow utility with respect to consumption:

$$\Lambda = e^{-\rho t}(C_t S_t)^{-\gamma} \quad (\text{IA.40})$$

Applying Itô's Lemma yields the SDE:

$$\begin{aligned} \frac{d\Lambda}{\Lambda} = & \left(-\rho - \gamma\mu_{ct} + \frac{\gamma^2}{2}\sigma_{ct}^2 \right) dt - \gamma\sigma_{ct} dW_{ct} \\ & + \underbrace{\left(-\gamma \log(\phi)x_t + 2\rho_{cx}\sigma_{ct}\sigma_{xt} + \sigma_{xt}^2 \right)}_{\text{habit model only}} dt - \gamma\sigma_{xt} dW_{xt} \end{aligned} \quad (\text{IA.41})$$

Recursive utility. Following Tsai and Wachter 2018, the value function takes the form:

$$V_t = \frac{C_t^{1-\gamma} e^{(1-\gamma)K(x_t)}}{1-\gamma} \quad (\text{IA.42})$$

where $K(x_t)$ captures the full dependence of the value function on the state variable (see Section IA.J.8 for the perturbation method to solve for K). The SDF is derived from the aggregator F as:

$$\begin{aligned} \frac{d\Lambda}{\Lambda} = & \left(\frac{\rho \left((1-\gamma\psi)e^{\frac{(1-\psi)K(x_t)}{\psi}} + \gamma\psi - \psi \right)}{1-\psi} - \gamma\mu_{ct} + \frac{\gamma^2\sigma_{ct}^2}{2} \right. \\ & + \frac{\gamma(\gamma\psi - 1)\rho_{cx}\sigma_{xt}\sigma_{ct}K'(x_t)}{\psi} \\ & + \frac{(\gamma\psi - 1) \left(2\psi(\mu_{x0} - x_t) \log(\phi)K'(x_t) + \sigma_{xt}^2 ((\gamma\psi - 1)K'(x_t)^2 - \psi K''(x_t)) \right)}{2\psi^2} \Bigg) dt \\ & - \frac{(\gamma\psi - 1)\sigma_{xt}K'(x_t)}{\psi} dW_{xt} - \gamma\sigma_{ct} dW_{ct} \end{aligned} \quad (\text{IA.43})$$

When $\gamma = 1/\psi$ this reduces to the TSU formula.

IA.J.5 Short-term interest rate

The short rate is $r(x_t) = -E_t[d\Lambda/\Lambda]/dt$. For TSU:

$$r(x_t) = \rho + \gamma\mu_{ct} - \frac{\gamma^2}{2}\sigma_{ct}^2 - \underbrace{\gamma \log(\phi)x_t + 2\rho_{cx}\sigma_{ct}\sigma_{xt} + \sigma_{xt}^2}_{\text{habit model only}} \quad (\text{IA.44})$$

For RU:

$$\begin{aligned}
r(x_t) = & \frac{\rho \left((1 - \gamma\psi) e^{\frac{(1-\psi)K[x_t]}{\psi}} + \gamma\psi - \psi \right)}{1 - \psi} + \gamma\mu_{ct} - \frac{\gamma^2\sigma_{ct}^2}{2} \\
& - \frac{\gamma(\gamma\psi - 1)\rho_{cx}\sigma_{xt}\sigma_{ct}K'(x_t)}{\psi} \\
& - \frac{(\gamma\psi - 1) \left(2\psi(\mu_{x0} - x_t) \log(\phi)K'(x_t) + \sigma_{xt}^2 ((\gamma\psi - 1)K'(x_t)^2 - \psi K''(x_t)) \right)}{2\psi^2}
\end{aligned} \tag{IA.45}$$

IA.J.6 Bond pricing

The price $Q(x_t, m)$ of a zero-coupon bond with remaining maturity m satisfies the PDE:

$$-Q_m - r(x_t)Q + \left(-\log(\phi)(\mu_{x0} - x_t) + A(x_t) \right) Q_x + \frac{\sigma_{xt}^2}{2} Q_{xx} = 0 \tag{IA.46}$$

where subscripts denote partial derivatives and $A(x_t) dt \equiv (d\Lambda/\Lambda)(dQ)$ is the covariance between the SDF and the bond price (i.e., the source of the term premium). The terminal condition is $Q(x_t, 0) = 1$.

The solution is given by the Feynman–Kac formula:

$$Q(m, x_t) = E_t \left[\exp \left\{ - \int_0^m r(\tilde{x}_{t+s}) ds \right\} \right] \tag{IA.47}$$

where $\tilde{x}$ follows the risk-adjusted process:

$$d\tilde{x}_t = \left(-\log(\phi)(\mu_{x0} - \tilde{x}_t) + A(\tilde{x}_t) \right) dt + \sigma_{xt}(\tilde{x}_t) dW_{xt} \tag{IA.48}$$

This expectation is computed via Monte Carlo simulation.

IA.J.7 Risk-neutral yield and term premium

Setting $A \equiv 0$ gives the risk-neutral bond price:

$$H(m, x_t) = \mathbb{E}_t \left[\exp \left\{ - \int_0^m r(x_{t+s}) \, ds \right\} \right] \quad (\text{IA.49})$$

The term premium of maturity m is:

$$TP(x_t, m) = \frac{-\log Q(x_t, m) - (-\log H(x_t, m))}{m} \quad (\text{IA.50})$$

IA.J.8 Perturbation approximation for K (recursive utility)

In the RU case, $K(x_t)$ is the solution of the following ordinary differential equation (ODE), obtained by substituting the value function into the Hamilton–Jacobi–Bellman equation with $\psi = 1/(1 - \epsilon)$:

$$\begin{aligned} -\frac{1}{2}\gamma\sigma_{ct}^2 + \frac{1}{2}\sigma_{xt}^2 \left(K''(x_t) - (\gamma - 1)K'(x_t)^2 \right) - \log(\phi)(\mu_{x0} - x_t)K'(x_t) \\ + \frac{\rho(e^{-\epsilon K(x_t)} - 1)}{\epsilon} + \mu_{ct} = 0 \end{aligned} \quad (\text{IA.51})$$

For $\psi = 1$ (i.e. $\epsilon = 0$) this ODE has an analytic solution $K_0(x_t)$. Expanding K in powers of ϵ :

$$K(x_t; \epsilon) = K_0(x_t) + \epsilon K_1(x_t) + \epsilon^2 K_2(x_t) + \dots \quad (\text{IA.52})$$

and substituting into (IA.51), each K_n is a polynomial in x_t of degree $n + 1$:

$$\begin{aligned}
K_0(x_t) &= a_{0,0} + a_{0,1}x_t \\
K_1(x_t) &= a_{1,0} + a_{1,1}x_t + a_{1,2}x_t^2 \\
K_2(x_t) &= a_{2,0} + a_{2,1}x_t + a_{2,2}x_t^2 + a_{2,3}x_t^3 \\
&\vdots
\end{aligned} \tag{IA.53}$$

The coefficients $a_{m,n}$ are determined by matching powers of $x_t^m \epsilon^n$ to zero, yielding a sequence of linear equations that can be solved successively. This provides a global approximation in the state variable for a wide range of IES values. The approach extends Tsai and Wachter 2018, who derived only the zeroth-order term K_0 ; higher-order terms enable accurate approximations well away from $\psi = 1$.³⁹

IA.K Model Variations: Yields and Term Premia

This section presents yield and term premium results for all 37 model variations. Each figure has four panels, with the respective state variable on the x-axis. The top-left panel shows the instantaneous rate, the five-year yield, and the five-year risk-neutral yield. The top-right panel shows the decomposition of the five-to-ten-year forward rate into the full yield, the risk-neutral yield, and the term premium. The bottom-left panel shows the instantaneous, five-year, and ten-year yields. The bottom-right panel shows the term premia for three-, five-, and ten-year bonds. The solid vertical line marks the ergodic median; dashed vertical lines show the 10th-percentile minimum and 90th-percentile maximum over 12-year simulations, so that 90% of simulated paths lie within the boundaries.

³⁹One limitation is that $\sigma_{x_t}^2$ and $\sigma_{c_t}^2$ must be at most linear in x_t , so unlike in the TSU case where $\sigma_{c_t} \propto x_t$, the RU case requires $\sigma_{c_t} \propto \sqrt{x_t}$, somewhat restricting the variability of consumption volatility.

IA.K.1 Model index

Abbreviations used: time-varying (TV), consumption drift (CD), consumption volatility (CV), time-separable utility (TSU), recursive utility (RU), intertemporal elasticity of substitution (IES).

Model Variation Description	Abbreviation	Reference
<i>Time-separable utility (TSU)</i>		
TV CD with TSU.	TSU-CD	Figure IA.5
TV CD with TSU and high risk aversion.	TSU-CD-HRA	Figure IA.6
TV CD with TSU and low persistence.	TSU-CD-LP	Figure IA.7
TV CD with TSU and high correlation ρ_{cx} .	TSU-CD-HCor	Figure IA.8
TV CD with TSU and high impatience.	TSU-CD-HImp	Figure IA.9
TV and high CD with TSU.	TSU-HCD	Figure IA.10
TV CD with TSU and high CV.	TSU-CD-HCV	Figure IA.11
TV CV with TSU.	TSU-CV	Figure IA.12
TV CV with TSU and high risk aversion.	TSU-CV-HRA	Figure IA.13
TV CV with TSU and high CD.	TSU-CV-HCD	Figure IA.14
TV and HCV with TSU and positive ρ_{cx} .	TSU-HCV	Figure IA.15
TV and HCV with TSU and negative ρ_{cx} .	TSU-HCV-NCor	Figure IA.16
Both TV CD and CV, rate <u>d</u> ecreasing in CV and ρ_{cx} <u>p</u> ositive.	TSU-Arb-DP	Figure IA.17
Both TV CD and CV, rate <u>i</u> ncreasing in CV and ρ_{cx} <u>n</u> egative.	TSU-Arb-IN	Figure IA.18
Both TV CD and CV, rate <u>d</u> ecreasing in CV and ρ_{cx} <u>n</u> egative.	TSU-Arb-DN	Figure IA.19
Both TV CD and CV, rate <u>i</u> ncreasing in CV and ρ_{cx} <u>p</u> ositive.	TSU-Arb-IP	Figure IA.20
TV external habit with TSU.	TSU-Habit	Figure IA.21
TV external habit with TSU and low b .	TSU-Habit-Low.b	Figure IA.22
TV external habit with TSU and $b < 0$.	TSU-Habit-Neg.b	Figure IA.23
TV external habit with TSU and constant state volatility.	TSU-Habit-CSV	Figure IA.24

Table IA.19: Index of model variations (Part I: time-separable utility).

Model Variation Description	Abbreviation	Reference
<i>Recursive utility (RU)</i>		
TV CD with RU.	RU-CD	Figure IA.25
TV CD with RU and high risk aversion.	RU-CD-HRA	Figure IA.26
TV CD with RU and high IES.	RU-CD-HIES	Figure IA.27
TV CD with RU and low IES.	RU-CD-LIES	Figure IA.28
TV CD with RU and high ρ_{cx} .	RU-CD-HCor	Figure IA.29
TV CD with RU and negative ρ_{cx} .	RU-CD-NCor	Figure IA.30
TV and high CD with RU.	RU-HCD	Figure IA.31
TV CD with RU and high CV.	RU-CD-HCV	Figure IA.32
TV and heteroskedastic CD with RU and ρ_{cx} positive.	RU-CD-Heterosk-PCor	Figure IA.33
TV and heteroskedastic CD with RU and ρ_{cx} negative.	RU-CD-Heterosk-NCor	Figure IA.34
TV CV with RU.	RU-CV	Figure IA.35
TV CV with RU and high risk aversion.	RU-CV-HRA	Figure IA.36
TV CV with RU and high persistence.	RU-CV-HP	Figure IA.37
TV CV with RU and high IES.	RU-CV-HIES	Figure IA.38
TV CV with RU and low IES.	RU-CV-LIES	Figure IA.39
TV and HCV with RU and ρ_{cx} positive.	RU-HCV-PCor	Figure IA.40
TV and HCV with RU and ρ_{cx} negative.	RU-HCV-NCor	Figure IA.41

Table IA.20: **Index of model variations (Part II: recursive utility).**

IA.K.2 TSU-CD (baseline calibration used in main article)




Figure IA.5: Time-varying CD with TSU. Top-left: instantaneous rate, 5-year yield, and 5-year risk-neutral yield as functions of consumption drift (or the respective state variable in other figures). Top-right: 5–10-year forward real yield decomposed into full yield, risk-neutral yield, and term premium. Bottom-left: instantaneous, 5-year, and 10-year yields. Bottom-right: term premia for 3-, 5-, and 10-year bonds. The solid vertical line marks the ergodic median; dashed vertical lines show the 10th-percentile minimum and 90th-percentile maximum over 12-year simulations.

IA.K.3 TSU-CD-HRA, $\gamma = 8$

Term premia are slightly larger but remain negative and constant with respect to the state variable.




Figure IA.6: Time-varying CD with TSU and higher risk aversion. See Figure IA.5 for plot details.

IA.K.4 TSU-CD-LP, $\phi = 0.8$

Term premia are unchanged. There is larger separation between yields, consistent with the expectation mechanism.




Figure IA.7: Time-varying CD with TSU and low persistence. See Figure IA.5 for plot details.

IA.K.5 TSU-CD-HCor, $\rho_{cx} = 1$

Term premia are larger in absolute value.




Figure IA.8: Time-varying CD with TSU and high correlation ρ_{cx} . See Figure IA.5 for plot details.

IA.K.6 TSU-CD-HImp, $\rho = 0.05$

Yields shift higher without any change in term premia.




Figure IA.9: Time-varying CD with TSU and high impatience. See Figure IA.5 for plot details.

IA.K.7 TSU-HCD, $\mu_{c0} = 0.06$

Yields shift higher without any change in term premia.




Figure IA.10: Time-varying and high CD with TSU. See Figure IA.5 for plot details.

IA.K.8 TSU-CD-HCV, $\sigma_{ct} = 0.16$

Yields shift down and term premia increase in absolute value but remain approximately constant.




Figure IA.11: Time-varying CD with TSU and high CV. See Figure IA.5 for plot details.

IA.K.9 TSU-CV (baseline calibration used in main article)




Figure IA.12: Time-varying CV with TSU. See Figure IA.5 for plot details.

IA.K.10 TSU-CV-HRA, $\gamma = 8$

Term premia increase in absolute value but not enough to turn positive, and yields become very high.




Figure IA.13: Time-varying CV with TSU and high risk aversion. See Figure IA.5 for plot details.

IA.K.11 TSU-CV-HCD, $\mu_{c0} = 0.08$

Term premia are unchanged but yields shift implausibly high.




Figure IA.14: Time-varying CV with TSU and high CD. See Figure IA.5 for plot details.

IA.K.12 TSU-HCV (baseline calibration used in main article)




Figure IA.15: Time-varying and HCV with TSU and positive ρ_{cx} . See Figure IA.5 for plot details.

IA.K.13 TSU-HCV-NCor, $\rho_{cx} < 0$




Figure IA.16: Time-varying HCV with TSU and negative ρ_{cx} . See Figure IA.5 for plot details.

IA.K.14 Arb-DP (baseline calibration used in main article)




Figure IA.17: Both time-varying CD and CV with TSU, short-term rate decreasing in CV and positive ρ_{cx} . See Figure IA.5 for plot details.

IA.K.15 Arb-IN (baseline calibration used in main article)




Figure IA.18: Both time-varying CD and CV with TSU, short-term rate increasing in CV and negative ρ_{cx} . See Figure IA.5 for plot details.

IA.K.16 Arb-DN




Figure IA.19: Both time-varying CD and CV with TSU, short-term rate decreasing in CV and negative ρ_{cx} . See Figure IA.5 for plot details.

IA.K.17 Arb-IP




Figure IA.20: Both time-varying CD and CV with TSU, short-term rate increasing in CV and positive ρ_{cx} . See Figure IA.5 for plot details.

IA.K.18 TSU-Habit (baseline calibration used in main article)




Figure IA.21: Time-varying external habit with TSU. See Figure IA.5 for plot details.

IA.K.19 TSU-Habit-Low.b, $b = 0.033$

Term premia are unchanged but yields become flatter.




Figure IA.22: Time-varying external habit with TSU and low b . See Figure IA.5 for plot details.

IA.K.20 TSU-Habit-Neg.b, $b = -0.033$

The short-term rate is now procyclical and term premia are negative.




Figure IA.23: Time-varying external habit with TSU and $b < 0$. See Figure IA.5 for plot details.

IA.K.21 TSU-Habit-CSV, $\sigma_{xt} = \lambda(0)\sigma_{c0}$

Term premia become constant, since consumption volatility is constant in the habit model and state-variable heteroskedasticity is the only remaining source of time variation.




Figure IA.24: Time-varying external habit with TSU and constant state variable volatility. See Figure IA.5 for plot details.

IA.K.22 RU-CD (baseline calibration used in main article)




Figure IA.25: Time-varying CD with RU. See Figure IA.5 for plot details.

IA.K.23 RU-CD-HRA, $\gamma = 6$

Term premia remain constant and negative but become significantly larger in absolute value. This suggests that time-varying risk aversion would produce time-varying term premia — a mechanism similar to the habit model.




Figure IA.26: Time-varying CD with RU and high risk aversion. See Figure IA.5 for plot details.

IA.K.24 RU-CD-HIES, $\psi = 1.43$

Term premia do not change significantly. The range of the short rate increases.




Figure IA.27: Time-varying CD with RU and high IES. See Figure IA.5 for plot details.

IA.K.25 RU-CD-LIES, $\psi = 0.83$

Term premia do not change significantly. Curiously the range of the short rate increases again.




Figure IA.28: Time-varying CD with RU and low IES. See Figure IA.5 for plot details.

IA.K.26 RU-CD-HCor, $\rho_{cx} = 1$

Term premia increase in absolute value but do not double in size as the correlation parameter does.




Figure IA.29: Time-varying CD with RU and high ρ_{cx} . See Figure IA.5 for plot details.

IA.K.27 RU-CD-NCor, $\rho_{cx} < 0$

Term premia increase but remain negative, as the dominant term in A does not involve ρ_{cx} .




Figure IA.30: Time-varying CD with RU and negative ρ_{cx} . See Figure IA.5 for plot details.

IA.K.28 RU-HCD, $\mu_{x0} = 0.05$

Term premia are unchanged; yields increase.




Figure IA.31: Time-varying and high CD with RU. See Figure IA.5 for plot details.

IA.K.29 RU-CD-HCV, $\sigma_{ct} = 0.08$

Term premia are unchanged; yields decrease.




Figure IA.32: Time-varying CD with RU and high CV. See Figure IA.5 for plot details.

IA.K.30 RU-CD-Heterosk-PCor

When the state variable is heteroskedastic, term premia become time-varying. They remain negative here.




Figure IA.33: Time-varying and heteroskedastic CD with RU and positive ρ_{cx} . See Figure IA.5 for plot details.

IA.K.31 RU-CD-Heterosk-NCor

Despite the sign change in ρ_{cx} , term premia remain negative because the dominant component of A does not contain ρ_{cx} .




Figure IA.34: Time-varying and heteroskedastic CD with RU and negative ρ_{cx} .

See Figure IA.5 for plot details.

IA.K.32 RU-CV (baseline calibration used in main article)




Figure IA.35: Time-varying CV with RU. See Figure IA.5 for plot details.

IA.K.33 RU-CV-HRA, $\gamma = 6$

Term premia have hardly moved.




Figure IA.36: Time-varying CV with RU and high risk aversion. See Figure IA.5 for plot details.

IA.K.34 RU-CV-HP, $\phi = 0.96$

Term premia have hardly moved; yields become slightly more variable.




Figure IA.37: Time-varying CV with RU and high persistence. See Figure IA.5 for plot details.

IA.K.35 RU-CV-HIES, $\psi = 1.43$

Term premia have hardly moved; yields become slightly more variable.




Figure IA.38: Time-varying CV with RU and high IES. See Figure IA.5 for plot details.

IA.K.36 RU-CV-LIES, $\psi = 0.77$

Term premia have hardly moved; yields become slightly more variable.




Figure IA.39: Time-varying CV with RU and low IES. See Figure IA.5 for plot details.

IA.K.37 RU-HCV-PCor, $\sigma_{c0} = 0.14$, $\rho_{cx} = 0.5$

Term premia are positive: the first component of A (containing ρ_{cx}) dominates due to the increase in consumption volatility. Nevertheless, term premia are smaller than in the corresponding TSU case because the second component of A remains negative.




Figure IA.40: Time-varying HCV with RU and positive ρ_{cx} . See Figure IA.5 for plot details.

IA.K.38 RU-HCV-NCor, $\sigma_{c0} = 0.14$, $\rho_{cx} = -0.5$

Both components of A are negative, so term premia are negative and larger in absolute value than in RU-CV.




Figure IA.41: Time-varying HCV with RU and negative ρ_{cx} . See Figure IA.5 for plot details.