A model of time-varying risk premia with habits and production*

Ian Dew-Becker

Duke University

February 20, 2014

Abstract

This paper builds on the production-based asset pricing literature to generate return predictability in general equilibrium. The model embeds habit formation into Epstein–Zin preferences to generate movements in risk aversion. In a production setting, it generates predictable equity returns and realistic consumption dynamics – R^2s in return forecasting regressions match empirical results at both short and long horizons, while consumption growth is unpredictable. Empirically, a novel model-implied forecast of equity returns yields an R^2 of 50 percent at the five-year horizon in post-war data. The model can be used to study the effects of time-varying risk premia on the macroeconomy.

1 Introduction

Stock prices are more volatile than can be explained by movements in expected dividends, while excess returns on the aggregate stock market are predictable over time. The two phenomena are connected: changes in the discount rates applied to future dividends can induce excess volatility in asset prices.[1] This paper develops a new preference specification with time-varying risk aversion that generates realistically predictable and volatile stock returns. Its contribution to the large production-based asset pricing literature is to improve the fit of the model to return forecasting

*I appreciate helpful comments from Jason Beeler, Eduardo Davila, David Laibson, Lars Lochstoer, Kelly Shue, Eric Swanson, Tomasz Strzalecki, Luis Viceira, seminar participants at Cornell, Harvard, MIT, the Federal Reserve Bank of New York, Stanford GSB, Haas, Sloan, Wharton, Stern, Olin, Booth, Fuqua, the Federal Reserve Bank of San Francisco, and the Texas Finance Festival, and, especially, my advisers, John Campbell, Effi Benmelech, and Emmanuel Farhi.

[1] Grossman and Shiller (1981); Leroy and Porter (1981); Shiller (1981).

results while maintaining previous successes in matching other features of the economy. While there are many models that can generate volatile risk premia when the consumption process is taken as exogenous, the innovation of this paper is to develop a habit-formation specification that delivers realistic results in a production setting with endogenous consumption choice.

A first goal of consumption-based asset pricing is to find a utility function that is consistent with the behavior of asset prices. That is, given a consumption process, a utility function should lead to high and volatile equity returns.[2] If we find such a utility function, the next question we ask is whether the consumption process we observe empirically is actually *consistent* with that utility function. That is, would agents with the proposed utility function choose to save and consume (and work and take leisure) in ways consistent with what we observe historically? So the second goal of consumption-based asset pricing is to find a utility function that endogenously generates both realistic asset returns *and* consumption dynamics. That second goal is one of the reasons for the existence of the production-based asset pricing literature.

There is now a large literature that studies the behavior of asset prices in general equilbrium models.[3] While that literature has had many successes, it has generally not focused on predictability in aggregate asset returns, even though empirical work finds that a large fraction of the variation in asset prices is driven by shifts in expected returns.[4] The goal of this paper is to develop a model that can explain the observed predictability of equity returns while preserving the successes of the previous literature in matching the basic behavior of the real economy, such as the volatilities of output, consumption, and investment growth.

The standard model of time-varying risk aversion is the habit specification of Campbell and Cochrane (1999).[5] In their model, when a person's consumption approaches her habit, her risk aversion rises. Using aggregate consumption data, Campbell and Cochrane find that their implied risk aversion measure can explain a large proportion of the movements in the price-dividend ratio on the stock market.

[2] This is essentially the question Mehra and Prescott (1985) ask.

[3] Particularly relevant to this paper are Jermann (1998), Lettau and Uhlig (2000), Tallarini (2000), Boldrin, Christiano, and Fisher (2001), Rudebusch and Swanson (2008), Guvenen (2009), Campanale, Casto, and Clementi (2010), De Graeve et al. (2010), and Gourio (2012), but there are many others.

[4] Gourio's (2012) model of time-varying disaster risk is the only published paper that generates R^2s in return forecasting regressions that are as large as we observe in the data.

[5] Other early papers studying habit formation include Abel (1990), Constantinides (1990), Jermann (1998), and Boldrin, Christiano, and Fisher (2001).

Campbell and Cochrane (1999) and other habit-formation models take an important step forward in that they are able to generate high and volatile risk premia (unlike power utility) under the assumption that consumption growth is i.i.d.. But in production settings, Campbell–Cochrane preferences imply that consumers strongly smooth consumption growth, leading to a consumption growth process that is far from i.i.d..[6] Moreover, since the consumption process with production is no longer what Campbell and Cochrane assumed it to be, risk premia in the production economy are not nearly as high or volatile as in their original endowment economy. Those problems are pervasive with models of habit formation, appearing also, for example, in Boldrin, Christiano, and Fisher (2001) and Jermann (1998).

The basic innovation of this paper is to embed the useful intuition in the previous literature—that persistent external habits can induce time-varying risk aversion—into generalized recursive preferences (Kreps and Porteus, 1978; Epstein and Zin, 1989; Weil, 1989). The Epstein–Zin specification permits the separate modeling of risk aversion and intertemporal substitution, while the Campbell–Cochrane intuition motivates time-variation in risk aversion. By separating variation in risk aversion from intertemporal substitution, the model resolves the usual problem of excess smoothness in consumption growth in models with strong habit formation.[7]

In the model, consumers have a time-varying external habit that is a benchmark to which they compare their own lifetime utility. That is, the habit is in the continuation value instead of current consumption. The higher is lifetime utility above the benchmark, the lower is risk aversion over proportional shocks to future welfare. I refer to the new preference specification as the EZ-habit model for its combination of these two frameworks. The basic mechanism, that bad news about wealth or future income raises risk aversion, is consistent with micro evidence from a range of studies discussed below on portfolio choice and measures of risk aversion from lab experiments.

The simple real business cycle (RBC) model with fixed labor supply provides a transparent laboratory in which to study the effects of time-variation in risk aversion on the macroeconomy in general equilibrium. I find that the dynamics of real variables and real interest rates under the EZ-habit specification are highly similar to a model with Epstein–Zin utility and constant relative

[6] See Lettau and Uhlig (2000) and Rudebusch and Swanson (2008). I also replicate their results in this paper.

[7] While it is obviously possible to reduce the smoothness of consumption growth by reducing the importance of the habit, that also reduces the size and volatility of the equity premium, bringing the model closer to simple power utility.

risk aversion, unlike Campbell–Cochrane preferences. The model can match both the short- and long-run variances of output, investment, and consumption growth.

In addition to matching macro moments, the EZ-habit preferences improve the fit of the RBC model to financial moments. Previous habit-based models designed to generate high or volatile risk premia tend to have implausibly volatile interest rates, a flaw not found here.[8] The reasonable behavior of interest rates is an important innovation: the EZ-habit model has stable interest rates but still generates substantial asset price volatility because it has variation in discount rates on risky assets that is driven by variation in risk aversion. Movements in discount rates imply that asset returns should be predictable, and extensive tests show that the degree of predictability in the model is similar to what is observed in the data.

Variation in risk aversion not only raises the volatility of asset returns, it also increases the equity premium by roughly one third on average. The reason is that when risk aversion is high, lifetime utility is low. Since the volatility of the pricing kernel under Epstein–Zin preferences depends on the volatility of lifetime utility, variation in risk aversion increases the volatility of the pricing kernel and hence risk premia.

There are numerous empirical methods of forecasting stock returns, but most are not based on equilibrium theories. For example, regressions of stock returns on price-dividend ratios are motivated simply by an identity that links the price-dividend ratio to future returns and dividend growth. Under the EZ-habit model, though, risk aversion can be measured directly. Using data on consumption and wealth, I construct an empirical estimate of risk aversion and show that it is a strong forecaster of aggregate stock returns: it outperforms the price-dividend ratio, Lettau and Ludvigson's (2001) measure of the consumption-wealth ratio, and Campbell and Cochrane's (1999) excess consumption ratio. Because it does not rely on an unobservable latent process to drive risk premia, this result differentiates the EZ-habit model from models of time-varying disaster risk.[9] I also provide novel evidence that shocks to productivity forecast future stock returns, which is consistent with the EZ-habit model but not time-varying disaster risk.

In addition to matching the predictability of returns, the model also matches forecasting results for consumption growth. Lettau and Ludvigson (2001) find little ability to forecast consumption

[8] See, e.g., Jermann (1998) and Boldrin, Christiano, and Fisher (2001).

[9] See, e.g., Gourio (2012), and Wachter (2010).

growth using their measure of the consumption-wealth ratio. Campbell and Shiller (1988) obtain similar results for dividend growth. As in the empirical data, it is essentially impossible to forecast consumption growth in the EZ-habit model using the consumption-wealth ratio or interest rates.

In addition to helping a simple RBC model match the basic features of the real economy and asset markets, the EZ-habit model is useful for giving a mechanism with which future papers can study the linkages between risk premia and the real economy. It provides a plauisible and quantitatively realistic description of variation in risk premia, which can in the future be studied, for example, in models with more complex models of innovation and production to further understand how shifts in risk premia over time affect the real economy.

The paper is organized as follows. Section 2 discusses the preference specification and lays out the economic environment. Section 3 calibrates a production economy and compares its behavior to the data. Section 4 tests the empirical implications of the model for return forecasting, and section 5 concludes.

2 The model

2.1 Preferences

For an agent with a constant elasticity of intertemporal substitution (EIS), Epstein–Zin (1989) utility can be expressed as

$$V_t = \left\{ (1 - \exp(-\beta)) C_t^{1-\rho} + \exp(-\beta) \left[G^{-1} \left(E_t \left[G \left(V_{t+1} \right) \right] \right) \right]^{1-\rho} \right\}^{1/(1-\rho)} \tag{1}$$

for a function G, where C_t is consumption and E_t is the expectation operator conditional on information available at date t.[10] The term $G^{-1}(E_t[G(V_{t+1})])$ is a certainty equivalent. When there is no uncertainty about V_{t+1}, $G^{-1}(E_t[G(V_{t+1})]) = V_{t+1}$. The usual choice for G (going back to Weil, 1989, and Epstein and Zin, 1991) is power utility,

$$G^{Power}(V_{t+1}) = V_{t+1}^{1-\alpha} \tag{2}$$

[10] The preferences can be further generalized to study alternative time aggregators, instead of the constant elasticity of substitution form.

Epstein and Zin (1989) show that risk aversion for an agent with preferences of the form (1) depends on the coefficient of relative risk aversion for G, while the EIS is equal to $1/\rho$.

Now consider a habit-formation utility function for G,

$$G^{Habit}(V_{t+1}; H_t) = (V_{t+1} - H_t)^{1-\alpha} \tag{3}$$

Certainty equivalent functions involving G^{Habit} are related to those using G^{Power} in the same way that usual habit specifications, for example, Constantinides (1991), are related to time-separable power utility. Rather than caring only about the absolute level of their continuation utility, G^{Habit} says that agents care about the spread between continuation utility and a benchmark H_t. Because the utility function adds a habit to Epstein–Zin, I refer to it as the EZ-habit specification.[11] I refer to the version of V_t using G^{Power} for the certainty equivalent as EZ-CRRA.

The coefficient of relative risk aversion for G^{Habit} is equal to $\alpha\frac{V_{t+1}}{V_{t+1}-H_t}$. As the spread between lifetime utility and the habit rises, the coefficient of relative risk aversion falls. Intuitively, when the continuation value falls close to its benchmark, proportional shocks to V_{t+1} loom much larger than when the agent has a cushion between his continuation value and H_t. If H_t varies slowly over time, then innovations to risk aversion depend on innovations to lifetime utility. Risk aversion in this model is thus dependent on news about the future, in particular news about future consumption growth (either its level or its higher moments). In this regard it is rather different from Campbell and Cochrane (1999) and other previous habit setups, where risk aversion depends only on current and past shocks to consumption growth.

G^{Habit} has three important drawbacks. First, if the support of the shocks to V_{t+1} is sufficiently wide, there is a non-zero probability that V_{t+1} will fall below H_t, leaving the certainty equivalent undefined.[12] To operationalize G^{Habit}, then, we need to augment the model to ensure that V_{t+1} cannot fall below the habit. Second, because G^{Habit} is not log-linear in V_{t+1}, obtaining simple analytic results with it is difficult or impossible. Third, also because G^{Habit} is not log-linear,

[11] Other papers, e.g., Rudebusch and Swanson (2010) and Yang (2008), incorporate consumption habits into Epstein–Zin preferences. That is, the $C_t^{1-\rho}$ term is replaced by $(C_t - X_t)^{1-\rho}$ where X_t is the habit. Rudebusch and Swanson (2008) show that in general equilibrium this does not lead to a time-varying Sharpe ratio because households endogenously smooth consumption to reduce their overall risk exposure.

[12] This issue also arises in other habit specifications. When models are solved with standard perturbation methods, the problem is simply ignored. I use a more precise global numerical solution technique that forces me to grapple with the problem.

standard arguments for the existence of a representative agent when agents have different wealth do not apply.[13] With G^{Habit}, we need to use a more trivial sort of aggregation in which every agent has the same level of wealth and faces identical shocks.

All of that said, the results in the remainder of the paper go through using G^{Habit} (as long as we take sufficient steps to ensure that V_{t+1} cannot fall below the habit). For the sake of both theoretical and numerical tractability, though, I replace G^{Habit} in the remainder of the paper with the convenient and tractable alternative

$$\begin{aligned} G_t^{TV}\left(V_{t+1}\right) &= V_{t+1}^{1-\alpha_t} \\ \alpha_t &= \alpha\frac{V_t^A}{V_t^A - H_t^A} \end{aligned} \tag{4}$$

where superscript A denotes an aggregate variable and TV stands for time-varying. $G_t^{TV(-1)}\left(E_t G_t^{TV}\left(V_{t+1}\right)\right)$ is a second-order approximation to $G^{habit(-1)}\left(E_t G^{Habit}\left(V_{t+1}; H_t\right)\right)$ around the non-stochastic version of the model.[14] Moreover, the appendix shows that in the continuous-time limit (i.e., under stochastic differential utility), preferences with G^{TV} are exactly equivalent to preferences using G^{Habit}, regardless of the process driving lifetime utility or the habit.[15] G^{TV} is locally equivalent to G^{Habit} in terms of risk preferences, but it solves the problems of integrability inside the certainty equivalent and the existence of a representative agent.

As in Campbell and Cochrane (1999), I assume that agents take the excess welfare ratio, $\frac{V_t^A}{V_t^A - H_t^A}$, and hence the coefficient of relative risk aversion, α_t, as external to their own decisions. The final step, then, is to specify a dynamic process for risk aversion. I assume a simple log-linear process, similar to Campbell and Cochrane (1999),

$$\alpha_{t+1} = \phi\alpha_t + (1-\phi)\,\bar{\alpha} + \lambda\left(\Delta v_{t+1}^A - E_t\Delta v_{t+1}^A\right) \tag{5}$$

where v_t^A is the log of V_t^A for the representative agent. Intuitively, when lifetime utility unex-

[13] A representative agent may exist, but her preferences need not actually look like the preferences of any particular agent. Ideally, if every agent has identical preferences, the representative agent will also have those preferences.

[14] More precisely, the second-order approximation also assumes no growth. Adding a constant growth rate μ to V would change the result to $\alpha_t = \alpha\frac{(1+\mu)V_t}{(1+\mu)V_t - H_t}$. The remainder of the analysis is identical.

[15] Melino and Yang (2003) study a utility function with the same form as G^{TV}, but they take α_t as a latent variable and give no theoretical motivation for its variation. This paper is original for proposing inserting habits into the certainty-equivalent part of Epstein–Zin preferences to motivate movements in α_t and studying the central role of endogenous consumption choice.

pectedly rises, it moves away from the habit and risk aversion falls, so $\lambda < 0$. Movements in the habit, and hence risk aversion, depend on aggregate welfare so that they are not affected by an individual agent's decisions. The AR(1) specification for risk aversion is approximately equivalent to a specification where $\log H_t$ is a geometrically weighted moving average of past values of v_t^A.

The key drawback of G^{TV} with (5) is that it eliminates the interesting cross-sectional implications of G^{Habit}. For example, under G^{Habit} if a person receives a windfall compared to her neighbors, she should become less risk averse. This paper is concerned with the aggregate implications of habit formation, so I use G^{TV} for the sake of tractability and aggregation, but G^{Habit} may be of independent interest in future work, and it would allow for richer cross-sectional tests of habit formation in the continuation value.

The appendix derives the marginal rate of intertemporal substitution (the stochastic discount factor, or SDF) for the general form of Epstein–Zin preferences in (1). In the case of G^{TV}, we end up with the expression,

$$M_{t+1} \equiv \frac{\partial V_t / \partial C_{t+1}}{\partial V_t / \partial C_t} = \exp(-\beta) \frac{V_{t+1}^{\rho-\alpha_t}}{\left(E_t V_{t+1}^{1-\alpha_t}\right)^{\frac{\rho-\alpha_t}{1-\alpha_t}}} \frac{C_{t+1}^{-\rho}}{C_t^{-\rho}} \tag{6}$$

with the only difference from the SDF under canonical Epstein–Zin preferences being the subscript on α_t. The SDF is a critical piece of the model because its volatility determines the price of risk in the economy.[16] As usual, changes in expected consumption growth or volatility will affect the SDF through their effects on V_{t+1}. Changes in α_{t+1} (or H_{t+1}) will also affect the SDF in the same way. Specifically, when the habit rises and agents are more risk averse, they penalize consumption uncertainty more, driving V_{t+1} down. High risk-aversion states thus have high state prices.

It is also straightforward to derive the standard result that

$$W_t = V_t^{1-\rho} C_t^{\rho} / (1 - \exp(-\beta)) \tag{7}$$

where W_t is the equilibrium price of a claim on the agent's consumption stream, which I call the aggregate wealth portfolio. This formula holds regardless of whether risk aversion varies over time,

[16] Hansen and Jagannathan (1991) show that the maximum Sharpe ratio (expected excess return divided by standard deviation) attained by any asset in the economy is equal to the standard deviation of the SDF divided by its mean.

and leads to the familiar result from Epstein and Zin (1991),

$$M_{t+1} = \exp(-\beta)^{\frac{1-\alpha_t}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\rho\frac{1-\alpha_t}{1-\rho}} R_{w,t+1}^{\frac{\rho-\alpha_t}{1-\rho}} \tag{8}$$

where $R_{w,t+1}$ is the return on the wealth portfolio.

2.2 Discussion and micro evidence

The EZ–habit model is an extension of standard habit-based preferences. Rather than consumers having a habit level of consumption that they target, it assumes they have a habit level of lifetime utility. Because equation (7) shows that there is a direct link between lifetime utility and wealth, we could also think of the model as saying that agents have a benchmark level of wealth. The house-money effect of Thaler and Johnson (1990) has a somewhat similar intuition. Through lab experiments, they find that subjects who have recently gained money in betting games play more aggressively, consistent with the EZ-habit model.[17]

Abel (1990) interprets habits in consumption as a "keeping up with the Joneses" effect. That intuition extends to the EZ-habit model. What agents try to keep up with in this model, though, is fundamentally different. For example, consider a college senior who is trying to decide between following her friends into consulting versus getting a law degree. With the J.D., she knows that in the short run her consumption will be lower than that of her friends, but in the long run she will likely be better off. In a model with an external consumption habit, three years of consumption below that of her friends looks painful. But when the habit appears as a function of lifetime utility, the student is comfortable giving up consumption in the short run as long as she knows she will do well compared to her friends in the long run. Because the habit appears only in the risk aggregator, an agent with EZ-habit preferences is willing to substitute consumption over time in a way that an agent with standard habit-forming preferences is not. For the same reason, the EZ-habit model is not inconsistent with the mixed evidence on the effects of classic consumption habits at the micro level (e.g., Dynan, 2000).

A number of papers use financial investment choices to measure variation in risk aversion.

[17] Barberis, Huang, and Santos (2001) embed the house-money effect in a full asset-pricing model. See Gertner (1993) and Post et al. (2008) for evidence on the house-money effect from game shows.

Carroll (2002) and Bucciol and Miniaci (2011) find that investors with higher wealth tend to tilt their portfolios towards more risky assets. While Brunnermeier and Nagel (2008) argue that inertia is the dominant feature of behavior in portfolio choice, Calvet, Campbell, and Sodini (2009), after controlling for inertia effects, find a strong and significant relationship between innovations to wealth and the riskiness of an investor's portfolio. Calvet and Sodini (2010) show that higher past income, controlling for current wealth and genetic differences in risk attitudes, is also negatively related to the share of household portfolios invested in risky assets. Last, Paravisini, Rappoport, and Ravina (2013) find that investors in Lending Club choose less risky portfolios following negative shocks to their local home prices. Overall, then, with the notable exception of Brunnermeier and Nagel (2008), the empirical literature supports the idea that increases in wealth reduce risk aversion.

In addition to the portfolio-choice literature, there are numerous papers that measure variation risk preferences more directly. Guiso et al. (2013) find in a sample of Italian bank customers that aversion to hypothetical gambles rose following the financial crisis. Similarly, Cohn et al. (2013) find that when financial managers are primed to think about past negative market outcomes, they tend to choose less risky portfolios in a lab setting (where they were able to choose to gamble roughly \$220 of their own money). Furthermore, related papers find that general attitudes towards risk, for example loss aversion or a desire for certainty, tend to rise following negative events, whether they are financial losses (Tanaka, Camerer, and Nguyen, 2010), natural disasters (Cameron and Shah, 2012), or violence in war (Callen et al., 2013).[18] These studies, along with those on portfolio choice, provide micro evidence consistent with the EZ-habit model that risk aversion rises following shocks that reduce wealth, and thus lifetime utility. These changes occur due to both aggregate and purely idiosyncratic events.

2.3 Production

Aggregate output is a function of the capital stock, K_t, and productivity A_t

$$Y_t = A_t^{1-\gamma} K_t^{\gamma} \tag{9}$$

[18] Cohn et al. (2013) list a number of other citations in this area.

The production function (9) can be thought of as Cobb–Douglas with labor supply held fixed at unity.

The aggregate resource constraint is

$$K_{t+1} = (1 - \delta)\, K_t + Y_t - C_t$$

where δ is the depreciation rate of capital.

For the benchmark calibration, productivity follows a random walk in logs,

$$\begin{aligned} \log A_{t+1} &= \log A_t + \mu + \sigma_a \varepsilon_{t+1} \qquad (10) \\ \varepsilon_{t+1} &\sim \mathcal{N}(0,1) \end{aligned}$$

The drawback of using random-walk technology is that it is difficult to generate the degree of volatility for output and investment that is observed in the data.[19] I therefore also consider a dual-shock version of the model that can match both the short- and long-run variances of output,

$$\begin{aligned} A_t &= \bar{A}_t X_t \qquad (11) \\ \log \bar{A}_{t+1} &= \log \bar{A}_t + \mu + \sigma_a \varepsilon_{t+1} \qquad (12) \\ \log X_{t+1} &= \phi_x \log X_t + \sigma_x \varepsilon_{x,t+1} \qquad (13) \\ \varepsilon_{t+1}, \varepsilon_{x,t+1} &\sim i.i.d.\ \mathcal{N}(0,1) \qquad (14) \end{aligned}$$

$\bar{A}_t$ here is the permanent component of output, while X_t can be interpreted as a device to capture forces that drive short-run fluctuations in output and consumption, such as shocks to monetary policy or energy prices. I refer to the version of the model with random-walk technology as the benchmark model, whereas the model with permanent and temporary technology shocks is the dual-shock model.[20]

[19] In particular, without a mean-reverting component, it is impossible for the model to replicate the result from Cochrane (1994) that the long-run variance of output is smaller than the unconditional variance.

[20] Note that X_t does *not* represent any type of long-run risk shock. In Bansal and Yaron's (2004) model, there is a shock that has persistent effects on growth rates. X_t here has a persistent effects on the *level* of productivity. It has no effect on the long-run variance of consumption growth.

3 Calibration and simulation

I solve the model with projection methods, which entails fitting a polynomial approximation to the decision rule and searching for coefficients so that the equilibrium conditions hold exactly at certain specified points in the state space (see Judd, 1999). The Euler equation errors in the simulations imply that agents misprice a claim on capital by uniformly less than 1/100th of one basis point (i.e., one part in one million) over the range of the state space that the simulations visit, and the median simulated error is an order of magnitude smaller.

The model is parameterized to match quarterly data. Table 1 lists the parameter values and the target moments. Many of the parameters, such as the exponent on capital in the production function, take standard values. I discuss here the parameters that are unique to this paper or do not have standard and agreed-upon values.

I set $\rho = 2/3$ as in Bansal and Yaron (2004), for an EIS of 1.5. Bansal and Yaron note that an EIS greater than 1 is necessary for increases in volatility to lower asset prices (specifically, the wealth-consumption ratio) in an endowment economy. Similarly, an EIS greater than 1 ensures that increases in risk aversion increase the expected return on the wealth portfolio and lower its current price. Many studies attempting to estimate the EIS have obtained values much smaller than 1 (Hall, 1988; Campbell and Mankiw, 1989). An important test of the model will be whether it can match the empirical regressions even though the calibrated EIS is larger than 1.

I choose the variance of permanent innovations to technology to match the long-run variance of consumption growth in the data. Since technology and consumption are cointegrated in the model, the long-run variance of consumption growth is equal to the variance of the permanent technology shocks. I estimate the empirical long-run variance of consumption growth with a third-order univariate AR model (where the lag length was selected with the Bayesian information criterion) and obtain a value of 0.0088^2. That is, the quarterly innovations to the permanent component of consumption have a standard deviation of 0.88 percent. For the dual-shock model, I select the parameters σ_x and ϕ_x to match the short-run volatility of consumption and output growth. The parameters imply that the temporary component of technology has an unconditional standard deviation of 2.7 percent.[21]

[21] Smets and Wouters (2007) estimate that the 1-quarter autocorrelation of stationary technology shocks is 0.95. On the other hand, the one-quarter autocorrelation of detrended real GDP is 0.85. I take $\phi_x = 0.90$ as the midpoint

The persistence of risk aversion, ϕ, is set to match the empirical persistence of the price-dividend ratio for the aggregate stock market, as in Campbell and Cochrane (1999). The mean and volatility of risk aversion ($\bar{\alpha}$ and, implicitly, λ) are chosen to match the average Sharpe ratio for the stock market in the postwar sample and the degree of predictability observed using the price-dividend ratio to forecast stock returns. Mean risk aversion is 14, and the standard deviation is set to 6.2.[22]

I simulate a broad range of models. The main comparisons are between the EZ-habit and the canonical Epstein–Zin model with constant relative risk aversion (EZ-CRRA). I also report results for Campbell–Cochrane (1999) preferences. Finally, I report results using the same basic production setup as above but with adjustment costs in investment and habit formation in preferences as in Jermann (1998). The appendix describes the details of the latter two models. The key differences between the preferences used by Campbell and Cochrane (1999) and Jermann (1998) is that the habit accounts for a much smaller fraction of consumption in Jermann's calibration than in Campbell and Cochrane's.

The remainder of this section reports an analysis of the behavior of the EZ-habit preferences. I begin with a comparison of basic moments across the models. I next examine in detail the degree of predictability implied by the model. Third, I study impulse response functions to help understand how the EZ-habit preferences change the dynamic response of the economy to shocks compared to standard Epstein–Zin preferences. Finally, the last section examines whether the EZ-habit model can match the empirical result that the EIS is usually estimated to be zero in aggregate regressions.

3.1 Comparisons across models

Table 2 reports basic moments from simulations of the five models along with corresponding values from the data.

3.1.1 Real variables

The first row shows that all of the models are calibrated to match the long-run variance of consumption exactly, which, under balanced growth, means that they also match the long-run variances of

between these two values.

[22] When $\alpha_t < 0$, I still use the standard Euler equation even though the household's optimization problem is convex. In the simulations, $\alpha_t < 0$ only 1.5 percent of the time. Treating households as if they are risk-neutral in periods when $\alpha_t < 0$ (i.e., censoring α_t at zero) has no discernible effect on the results.

output and investment growth.

Rows 2 through 4 give the standard deviations of quarterly output, consumption, and investment growth. Both the EZ-CRRA and single-shock EZ-habit models have volatilities for output and investment growth that are well below the empirical values. As designed, the dual-shock model rectifies this problem, matching both the short-run and long-run variances well. Both versions of the EZ-habit model match the empirical variance of consumption growth.

For the real variables, the Jermann model displays behavior similar to that of the other models. Campbell–Cochrane, however, is a notable outlier in terms of the standard deviation of consumption growth. Column 5 shows, as observed by Lettau and Uhlig (2000), that Campbell–Cochrane habits lead to much smoother consumption paths than what we obtain either under the weaker habits of Jermann or the EZ-habit model.

To summarize, rows 1 through 4 show that the EZ-habit model can capture the basic unconditional moments of output, consumption, and investment. Unlike with Campbell–Cochrane preferences, endogenous consumption growth is not excessively smooth – in fact, it is more volatile than the standard Epstein–Zin model.

3.1.2 Financial variables

Rows 5 through 9 of table 2 summarize the financial side of the model and show that the EZ-habit model improves substantially on the EZ-CRRA, Campbell–Cochrane, and Jermann setups.

Row 5 reports the Sharpe ratio on an equity portfolio. Equity is modeled in the simulations as a levered claim on aggregate dividends. The leverage ratio is set to 2.4, which matches the cointegrating parameter between consumption and dividends in the U.S. data.[23]

Row 5 shows that the EZ-habit model generates a realistic Sharpe ratio of 0.32, and one that is higher by half than the 0.22 that we find with EZ-CRRA preferences. The higher Sharpe ratio is due to the fact that bad states – when technology growth is low – are also high risk aversion states in the EZ-habit model. Those states are thus doubly bad for lifetime utility: consumption is low and the agent is more averse to the future uncertainty she faces. That increased variability in lifetime utility leads to more variable state prices and hence higher Sharpe ratios (through a higher

[23] Specifically, if dividends are denoted D_t, the levered dividend is D_t^{λ}, where λ is the leverage ratio. The value of λ I use is similar though somewhat smaller than the leverage ratio of 2.74 used for a consumption claim by Abel (1999) and Gourio (2012), among others. Dividends are defined as total payments to capital.

Hansen–Jagannathan (1991) bound). The EZ-habit model can thus generate substantially higher risk premia for a given level of risk aversion than the EZ-CRRA model does.

A surprising result in row 5 is that the Sharpe ratio is only 0.07 for Campbell–Cochrane. That model was designed to generate a high and volatile risk premium. But because consumers endogenously choose such a smooth consumption path, the actual quantity of risk in the economy is small. Specifically, in models based on power utility with a consumption habit, the pricing kernel is,

$$M_{t+1} = \beta \left(\frac{C_{t+1} - H_t}{C_t - H_{t-1}} \right)^{-\alpha} \tag{15}$$

In Campbell and Cochrane's (1999) calibration (also used here), the habit is a very large fraction of consumption on average (94.3 percent) so the effective coefficient of relative risk aversion is also large. However, consumers endogenously choose a very smooth consumption path – a standard result under habit formation.

There are thus two opposing effects: a large habit raises risk aversion, but it also drives agents to endogenously smooth consumption growth. The latter effect turns out to dominate quantitatively. With highly smooth consumption growth, M_{t+1} has low volatility in the simulations, leading to low risk premia (the same result previously obtained by Lettau and Uhlig, 2000, and Rudebusch and Swanson, 2008). Jermann (1998) solves the problem of excess smoothness in consumption growth by adding investment adjustment costs, but we will see below that that solution has the fatal drawback that it leads to extremely volatile interest rates; in that model consumers demand very high rates of return in order to accept a consumption process as volatile as we observe in the data.

The EZ-habit model avoids the drawbacks of both the Campbell–Cochrane and Jermann setups by separating risk aversion from intertemporal substitution. Unlike with models based on power utility, shifts in risk aversion in the EZ-habit setup do not affect intertemporal substitution. Moreover, because the pricing kernel is driven by shocks to both current and future consumption, whether consumers endogenously smooth consumption growth or not has little effect on the pricing kernel.

Rows 6 and 7 next report the size of the equity premium in the models. While the EZ-habit model is clearly able to generate a high *price* of risk, the question is whether there is enough actual risk in equities to generate a realistically high equity premium. The two EZ-habit models generate means and volatilities for returns that are far closer to the equity return observed in the data than the EZ-CRRA model does. The mean equity premium rises from 0.9 to 4.1 percent per year (compared to 6.8 percent in the data), while the standard deviation rises from 4.1 to 10.6 percent.

The intuition for that result is straightforward. After a positive technology shock, not only do dividends rise, but discount rates fall since risk aversion falls. Prices thus rise through both cash-flow and discount-rate effects, and the returns on the wealth portfolio and the levered dividend claim are more volatile than under constant relative risk aversion, where there is only a cash-flow effect.[24]

That said, the equity premium in the EZ-habit model is still 270 basis points smaller than in the data. This paper does not include adjustment costs in investment for the sake of simplicity, but we can see that in column 6, Jermann's specification with adjustment costs does in fact yield equity returns that are slightly more volatile than in the EZ-habit model (though only by 1 percentage point).

Finally, rows 8 and 9 show that the EZ-habit model is able to generate reasonable behavior for the risk-free rate. Specifically, it is low and stable, as in the data.[25]As mentioned above, the Jermann model generates interest rates that are too volatile by an order of magnitude – the standard deviation in the Jermann model is 10.20 percent, compared to 1.16 percent in the data (Boldrin, Christiano, and Fisher, 2001, obtain a similar result).

3.1.3 Measures of predictability

The final three rows report measures of the predictability of consumption growth and excess equity returns in the models. First, row 10 reports the mean small-sample correlation of 20-quarter excess equity returns with the price/dividend ratio. The reported means are from 5000 simulations of 228-quarter samples. Under constant relative risk aversion, where the equity premium is nearly constant, the *population* correlation is roughly zero. Small-sample bias, though, leads to -0.25 in

[24] LeRoy and Porter (1981) and Shiller (1981) argue that dividends do not seem sufficiently volatile to explain the volatility of stock prices. Grossman and Shiller (1981) suggest that variation in discount rates can explain this puzzle.

[25] It is worth noting that the EZ-habit model generates a downward-sloping real term structure. This is consistent with empirical findings by Evans (1998). Furthermore, Dew-Becker (2013) finds that when the EZ-habit preferences are combined with a New-Keynesian model of the economy, they have no trouble matching the nominal term structure of interest rates.

table 2. The EZ-habit model strengthens the correlation to -45 percent and yields a degree of predictability that is nearly identical to what we observe in the data. The degree of predictability in the Campbell–Cochrane and Jermann models is quantitatively similar to that of CRRA and half as large as it is in the EZ-habit model. We thus see, as expected, that time-variation in risk aversion substantially increases the correlation between price/dividend ratios and future returns.

Row 11 reports the small-sample relationship between price/dividend ratios and future consumption growth (calculated in the same way as in row 10). Under the EZ-habit model, the predictive power of price/dividend ratios for future consumption growth is essentially zero, as we observe in the data. Both the EZ-CRRA and Campbell–Cochrane models, on the other hand, imply a large degree of predictability, with correlations 0.2 and 0.5, respectively, strongly at odds with the data.

Finally, row 12 shows that we obtain a similar result when we look at the correlation between interest rates and one-quarter-ahead consumption growth, which is one of the key implications of the EZ-habit model. Because risk aversion varies over time, there is a time-varying precautionary-saving effect (similar to the effect coming from stochastic volatility in Bansal and Yaron, 2004). The movements in the precautionary saving effect drive the correlation between interest rates and future consumption growth to nearly zero, as observed in the data. The other models in table 2, Campbell–Cochrane and Jermann in particular, imply very strong forecasting power for interest rates on one-quarter-ahead consumption growth, with correlations of 0.75 and -0.61, respectively.

To summarize, table 2 shows that the EZ-habit model can match a broad array of features of the economy: the short- and long-run variances of output, consumption, and investment growth; the means and standard deviations of the risk-free rate and Sharpe ratio on equities; and finally three basic measures of the predictability of equity returns and consumption growth. None of the other models examined here are able to match all of those features of the economy simultaneously.

3.2 Predictability in the simulated model

The primary reason to include time-varying risk aversion in the model is to generate predictability in equity returns. While table 2 reports some simple measures of the degree of predictability, we now study more closely how well the model can match observed predictability patterns at short and long horizons.

To give a benchmark for the amount of predictability observed empirically, table 1 plots R^2s from four separate univariate regressions of excess returns for the CRSP value-weighted index over various horizons on different predictive variables: the log price-dividend ratio (e.g., Campbell and Shiller, 1988, among many others); Lettau and Ludvigson's (2001) measure of the consumption-wealth ratio, *cay*; Campbell and Cochrane's (1999) excess consumption ratio; and an estimate of risk aversion derived from the EZ-habit model in section 4. For all four predictive variables, the R^2s generally rise as the sample length grows, and estimated risk aversion outperforms *cay*, excess consumption, and the price-dividend ratio.

The gray line labeled "Simulated mean" gives the mean R^2 from 5,000 regressions of excess equity returns on the price/dividend ratio over 228-quarter spans in the benchmark simulation of the single-shock model (the same length as the empirical sample).[26] The upper gray line gives the 95th percentile of the simulations. As in the data, the simulated R^2s rise as the horizon lengthens, from 6 percent at the 1-year horizon to 31 percent at 10 years. The model compares favorably with the empirical results for the price-dividend and excess-consumption ratios, with the simulated mean tracking the empirical values closely. The empirical R^2s for *cay* are higher, but still below the 95th percentile of the simulations. The only variable that the simulations cannot match is estimated risk aversion, but raising the volatility of risk aversion in the calibration would solve this problem.

The R^2s generated here are substantially higher than those obtained in production models such as Campanale, Castro, and Clementi's (2010) model of time-varying first-order risk aversion and Guvenen (2009) and De Graeve et al.'s (2010) studies of limited participation. Those papers obtain R^2s of 0.15 or less at the 5-year horizon, compared to 0.23 on average at the 5-year horizon for the EZ-habit model. The population R^2s are also essentially identical to those found by Wachter (2010) and Gourio (2012) in endowment-economy and production-based models, respectively, with time-varying disaster risk. Beyond them, no other models are able to generate this degree of predictability with endogenous consumption growth in a production setting.

It is clear from table 2 and figure 1, then, that the EZ-habit model is able to both generate a high price of risk (50 percent higher than Epstein–Zin with constant risk aversion) and a highly volatile price of risk. R^2s from simulated forecasting regressions are consistent with what we observe empirically, and help generate large movements in asset valuations. Moreover, table 2

[26] The median of the simulations is highly similar to the mean.

shows that while consumption growth is strongly predictable in the various other models studied, it is essentially unpredictable in the EZ-habit model, as in the data.

3.3 Impulse response functions

To help understand the dynamics of the model, figure 2 plots impulse response functions (IRFs) in the EZ-CRRA and benchmark EZ-habit models for four variables: consumption, lifetime utility, the risk-free rate, and the Sharpe ratio on the consumption claim. The shock is a unit standard deviation (88 basis-point) permanent increase in the level of technology, which leads to an identical long-run increase in consumption, capital, and output.

The top-left panel shows the response of lifetime utility. For the EZ-CRRA model, lifetime utility immediately jumps to a point just below its new steady state, and then slowly rises as agents accumulate capital. For the EZ-habit model, though, lifetime utility actually overshoots its new steady state. The reason is that the positive shock to productivity drives risk aversion down. When agents are less risk-averse, they place a higher value on their future consumption stream because they penalize uncertainty less strongly.

The overshooting of lifetime utility in figure 2 helps increase the volatility of the SDF (equation 6), raising the Hansen–Jagannathan bound. The top-right panel shows that on the impact of a shock, the Sharpe ratio in the EZ-habit model falls by 12.5 percent (as a fraction of its mean), and then gradually rises again, with a half-life of 12 quarters. That is the effect that generates the high degree of predictability in figure 1. The persistence of risk aversion is why the perdictability appears stronger at longer horizons.

The bottom-left panel shows the dynamics of the risk-free rate. The initial response is nearly identical for the two models. Since the size of the capital stock is essentially fixed in the short-run, an increase in productivity directly increases the return on capital. The agent's Euler equation links the expected return on capital to discount rates, so since the change in the return on capital is the same in the two models, the change in interest rates is also. Unlevered capital itself is a relatively low risk investment, so movements in risk aversion have little effect on capital discount rates. If capital were more risky, then we would expect the decline in risk aversion to lead to a relatively larger increase in investment and a smaller increase in consumption following the shock.

The final panel of figure 2 shows the response of consumption in the two models. The EZ-habit model shows a larger initial response of consumption, with lower expected consumption growth going forward. To see why this is, we can write consumption growth in a log-linearized version of the model as

$$E_t \Delta c_{t+1} = \bar{c} + \rho^{-1} r_{f,t+1} + \alpha_t \times vol \tag{16}$$

where $r_{f,t+1}$ is the risk-free interest rate between dates t and $t+1$, vol represents a measure of the total volatility in the model (which is fixed), and $\bar{c}$ is a constant. $\alpha_t \times vol$ represents the precautionary saving effect and is a function of the current level of risk aversion and the variances of the shocks in the model.

The standard interpretation in an endowment economy is that conditional on consumption growth, a strong precautionary saving motive leads to a low risk-free rate. In a production setting, though, it is the risk-free rate that is held roughly fixed because it is tied to the marginal product of capital, which is hard to change quickly through investment. Conditional on the risk-free rate, then, a small precautionary saving motive leads to lower expected consumption growth (more consumption today, saving less for tomorrow). In the EZ-habit model, a positive technology shock lowers risk aversion, and hence consumption rises more than in the canonical Epstein–Zin case. This effect also serves to increase the volatility of the SDF, just as the higher response of lifetime utility does.

Given the results in figure 2, it is straightforward to see what would happen in this economy if there were a pure shock to the coefficient of relative risk aversion. Since the risk-free rate is tied to the marginal product of capital, it would not move on the impact of a shock. The only effect on real variables of a pure decline in risk aversion, then, is that agents would want a smaller buffer stock of savings, so they would raise consumption and lower investment: shocks to risk aversion look like simple consumption demand shocks.[27]

3.4 Estimating the EIS from interest rate regressions

The value of the elasticity of intertemporal substitution is controversial. Regressions based on aggregate consumption and asset returns often find a very small EIS (Hall, 1988; Campbell and

[27] If unlevered capital were sufficiently risky, it is possible that a decline in risk aversion would actually raise rather than lower investment. However, consumption and investment would always be driven in opposite directions. In a New-Keynesian model with more frictions, however, they can be driven in the same direction.

Mankiw, 1989).[28] This result is in conflict with the calibration used here and in other recent production-based asset pricing studies (Kaltenbrunner and Lochstoer, 2010; Gourio, 2012), which assume that the EIS is greater than 1. The question is whether the EZ-habit model generates small EIS estimates in regressions similar to those estimated in Campbell (2003) and elsewhere. Table 3 shows that the EZ-habit model can in fact match the results of empirical EIS regressions.

Standard aggregate EIS regressions start from a model in which the risk-free rate takes the form

$$E_t \Delta c_{t+1} = b_0 + \rho^{-1} r_{f,t+1} \tag{17}$$

where ρ^{-1} is the EIS and b_0 is a parameter depending on the discount rate and underlying volatility in the model (as above). This relation is straightforward to derive in an endowment economy with homoskedastic consumption growth and in which agents have a constant EIS and coefficient of relative risk aversion. It is also obtained in a log-linearization of the standard RBC model with homoskedastic technology shocks. Note that equation (17) does not hold in the EZ-habit model; equation (16) is the true equation (up to a log-linearization).

In the simulations of the model in section 3, we have the ability to directly measure $E_t \Delta c_{t+1}$. The first row of table 3 reports the population estimate of ρ^{-1} in regression (17) using the true expectation of consumption growth under the EZ-CRRA and EZ-habit models. In the EZ-CRRA case, the regression identifies ρ^{-1} exactly, as it should. On the other hand, the estimate of ρ^{-1} is only 0.64 in the benchmark EZ-habit model, far closer to the value of roughly zero in the data.

The bias comes from the fact that the time-varying precautionary saving effect from equation (16) is omitted from the regression. Because precautionary saving is correlated with both expected consumption growth and interest rates, omitting it biases the usual EIS regression (17). An alternative way to see the source of the bias is to go back to the IRFs in figure 2. In both models the risk-free rate rises by the same amount following a shock. In the EZ-habit specification, though, because of the decline in precautionary saving, expected consumption growth is lower following a

[28] Campbell (2003) reviews the literature and estimates the EIS using a variety of specifications and data from a broad range of countries, finding values generally less than 0.5 and often less than 0.2. Vissing-Jorgenson (2002) finds an EIS less than unity in micro data. On the other hand, Vissing-Jorgenson and Attanasio (2003) and Gruber (2006) obtain larger estimates using micro data, both above unity. Gruber (2006) is particularly well-identified, using variation in the capital income tax rate as the source of exogenous differences in the after-tax interest rate earned by households.

shock than in the EZ-CRRA case. That means that the estimate of ρ^{-1} must fall.[29]

The regression in the first row of table 3 is in some sense ideal, but it is not the regression that we are actually able to run in the data since $E_t \Delta c_{t+1}$ is unobservable.[30] Rows 2 through 4 report results for estimates of ρ^{-1} from regressions of actual consumption growth, Δc_{t+1}, on the risk-free rate, $r_{f,t+1}$. Row 2 gives the population estimates, while rows 3 and 4 give the median and 95-percent range of the estimates from 228-quarter simulations. With constant relative risk aversion, the population regression in row 2 again estimates the EIS exactly. The median estimate from the small-sample regressions in row 3 is 1.16. While the 95-percent range is wide, it only barely contains the point estimate from the data. So it is in principle possible for the EZ-CRRA model to generate an estimate of the EIS as small as what we observe in the data, but the probability is small (less than 10 percent).

In the EZ-habit models, the bias in EIS regressions is far larger. The population estimate of the EIS in the single-shock case is 0.56, and the median small-sample estimate is 0.03 – almost exactly what we observe empirically. For the dual-shock model, the estimates are only slightly better: the small-sample median estimate is 0.35. The EZ-habit model is thus able to closely match the empirical fact that EIS estimated from aggregate regressions is near zero.

Given that in the model the bias in the EIS regressions comes from movements in risk aversion, if we could observe α_t we could completely eliminate the bias. The final three rows of table 3 try to estimate the EIS including a control for risk aversion.

In the data, I use a measure of risk aversion derived from the EZ-habit model below in section 4, denoted $\hat{\alpha}_t$. The empirical estimate of the EIS is essentially unchanged from when $\hat{\alpha}_t$ is not included. In population, when α_t is included in the simulated regressions, the EIS is estimated exactly, as expected. In small-sample regressions, though, the estimate of the EIS in the model is still biased downward. In the single-shock model, the median estimate is 0.07, and in the dual-shock model 0.94. Row 7 shows, however, that the 2.5 percentile of the small-sample estimates is -3.08 in

[29] In Bansal and Yaron (2004), time-variation in the volatility of shocks in principle causes EIS regressions to be biased. However, Beeler and Campbell (2012) show that their calibration generates almost no actual bias—the median sample EIS estimates are well above 1. This paper thus represents an improvement in being able to generate a substantial bias in aggregate regressions without large movements in the conditional volatility of consumption.

[30] In principle, the real risk-free rate, $r_{f,t+1}$, is also unobservable in the data. As above, I form $r_{f,t+1}$ as the difference between the nominal three-month interest rate and a forecast of inflation based on lagged inflation and nominal interest rates. Errors in the estimate of the true real-risk-free rate would bias the estimate of ρ^{-1} towards zero. In theory, instrumental-variables methods can eliminate this bias.

the dual-shock model, while the 97.5 percentile is 3.11. So even though the median estimate in the dual-shock model is not enormously biased, the empirical value of 0.18 is well within the simulated range. The reason for this result is that consumption follows nearly a random walk in the model, so there is very little variation from which to identify the EIS. In the end, although controlling for risk aversion should, in principle, allow us to estimate the EIS consistently, in small samples the regressions still do not seem to provide useful estimates because of weak-identification problems.

To summarize, the EZ-habit model is robustly able to match the empirical facts that estimates of the EIS from aggregate regressions of consumption growth on interest rates are near zero and that consumption growth is essentially unpreditable.

4 Empirical return forecasting

This section shows that under the EZ-habit model we can estimate risk aversion empirically up to an affine transformation, and that the estimated risk aversion proxy yields a strong forecast of stock returns. In fact, the forecast is stronger than all the other empirical predictors I examine. Moreover, I present novel evidence that shocks to TFP forecast future stock returns, as predicted by the model.

The forecasting results differentiate the EZ-habit model from models with time-varying disaster risk. Gourio's (2012) model predicts that when the probability that there will be a disaster changes, price-dividend ratios will forecast returns, which is also true in the EZ-habit model. But the EZ-habit model also predicts that technology and estimated risk aversion will forecast stock returns, and that estimated risk aversion will be the single most powerful forecaster of returns, which would not be true in the time-varying disaster model or models based on other forms of time-varying volatility. The evidence that estimated risk aversion and TFP shocks forecast returns thus argues in favor of time-varying risk aversion over time variation in disaster risk.

4.1 Estimating risk aversion

Given the AR(1) process for risk aversion in (5), it is straightforward to measure risk aversion if we simply observe the history of aggregate lifetime utility, v_t^A. For a given value of the EIS and observed data on wealth and consumption, it is possible to calculate v_t^A by rearranging equation

(7)

$$v_t^A = \frac{1}{1-\rho} w_t^A - \frac{\rho}{1-\rho} c_t^A + \frac{1}{1-\rho} \log\left(1 - \exp\left(-\beta\right)\right) \tag{18}$$

That is, the lifetime utility of the representative agent depends on aggregate wealth and aggregate consumption; if we can measure wealth and consumption, then we can measure lifetime utility. We then simply plug the estimates of v_t^A into the recursion for risk aversion (equation 5) to obtain estimates of α_t.

The primary difficulty with estimating aggregate wealth is that human wealth is not directly observable. I follow Lettau and Ludvigson (2001) in using labor income as a proxy for human wealth. Lettau and Ludvigson (2001) study a cointegrating relationhip between consumption, financial wealth, and labor income. Although their analysis was designed to estimate the consumption-wealth ratio, it also delivers, as a byproduct, a measure of aggregate wealth. I use that wealth measure to construct the measure of v_t^A and hence risk aversion.

Assuming the price-dividend ratio for human wealth is stationary, labor income is a proxy for human wealth. Denoting asset wealth as a_t and labor income as y_t, the appendix shows that we then have a cointegrating relation (the one used by Lettau and Ludvigson, 2001),

$$c_t = \zeta \omega a_t + \zeta \left(1 - \omega\right) y_t + \xi_t' \tag{19}$$

where ζ and ω are parameters and ξ_t' is a stationary error term. That is, consumption, asset wealth, and labor income are jointly cointegrated (a result which holds in the production model of the previous section). Lettau and Ludvigson (2001) refer to the residual ξ_t' as *cay*. This variable essentially represents an estimate of the consumption-wealth ratio. To measure wealth, then, we simply define

$$ay_t \equiv \omega a_t + \left(1 - \omega\right) y_t \tag{20}$$

which, under the assumptions above, will be a statistically unbiased estimate of total wealth. So our measure of aggregate wealth for use in equation (18) is a linear combination of asset wealth and labor income, where the relative weights are obtained from the estimated cointegrating relationship.

With our measure of wealth ay_t, we estimate v_t^A as

$$\hat{v}_t^A = \frac{1}{1-\rho} ay_t - \frac{\rho}{1-\rho} c_t \tag{21}$$

where we ignore the constant, and a circumflex indicates an estimated variable. To the extent that there is measurement error in the consumption or wealth data, $\hat{v}_t^A$ will inherit that same error. When we use $\hat{v}_t^A$ to forecast market returns, this measurement error should only weaken the results. For measurement error to generate a spurious predictive relation, it would have to be correlated with other predictors of returns.[31]

This definition of $\hat{v}_t^A$ is similar to Lettau and Ludvigson's cay_t, except they have equal weights on c_t and ay_t, whereas equation (21) uses a combination in which the weights depend on the EIS. Also, cay_t is stationary by construction, whereas $\hat{v}_t^A$ is growing over time (it is cointegrated with consumption and wealth).

In equation (21) a high EIS (low ρ) raises the weight on consumption relative to asset wealth. If the EIS is less than 1 ($\rho > 1$), the weight on wealth, ay_t, is actually negative, and the weight on consumption is greater than 1. Bansal and Yaron (2004) and Kaltenbrunner and Lochstoer (2010) both find that an EIS of 1.5 allows their models to fit asset pricing facts, so I use the same value. This value is also consistent with the micro evidence of Vissing-Jorgensen and Attanasio (2003). The results reported below are quantitatively similar as long as the EIS is greater than 1.1 (at that level and below, $\hat{v}_t^A$ becomes very volatile). The appendix reports a sensitivity analysis for various values of the EIS.

Finally, I construct an estimate $\hat{\alpha}_t$ using the update process for risk aversion, equation (5), and the data on $\hat{v}_t^A$,

$$\hat{\alpha}_{t+1} = \phi \hat{\alpha}_t + (1-\phi)\bar{\alpha} + \lambda \left(\Delta \hat{v}_{t+1}^A - E_t \Delta v_{t+1}^A\right) \tag{22}$$

As above, I assume that $\phi = 0.96$. $E_t \Delta v_{t+1}^A$ is estimated simply as the sample average of $\Delta \hat{v}_t^A$.[32]

[31] One obvious source of measurement error is that human capital is not a perfect estimator of the value of human wealth. Suppose risk aversion rises above average and lowers the price-dividend ratio on human wealth below average. Labor income will then be overestimating human wealth (compared to its average). High levels of wealth drive our measure of $\hat{\alpha}_t$ downward, so this measurement error should bias the results against correctly forecasting returns (high risk aversion in the data leads to low risk aversion in our estimates). In simulations not reported here, though, this effect is inconsequential.

[32] In principle, it is possible to forecast Δv_{t+1}^A, but the amount of predictability in Δv_{t+1}^A is sufficiently small that the results are nearly identical to assuming that v_{t+1}^A simply follows a random walk. The appendix also shows that

The parameter λ governs the volatility of α_t, but it has only a multiplicative effect on $\hat{\alpha}_t$. That is, any two estimates of $\hat{\alpha}_t$ will be perfectly correlated, regardless of what values are chosen for λ. The same argument applies for $\bar{\alpha}$. As long as we are simply trying to forecast stock returns using a linear regression, we can ignore any additive or multiplicative shifts in $\hat{\alpha}_t$. Therefore, I set $\bar{\alpha} = 0$ and choose λ so that $\hat{\alpha}_t$ has unit variance, normalizations that will not affect the regression-based measures of forecasting power (and I choose a negative value of λ to match the habit-formation motivation of the model). In the first period of the sample I assume $\hat{\alpha} = \bar{\alpha}$.

An important feature of this method of forecasting is that it is based only on the preference specification. No assumptions about the production side of the economy are required for this method to be valid. We simply take advantage of the relation between lifetime utility and changes in risk aversion and the relation under Epstein–Zin preferences between lifetime utility and wealth.

4.2 Forecasting market returns

The next question is to what extent the model-implied variation in expected returns is related to actual returns. Figure 3a plots $\hat{\alpha}_t$ and five-year forward-looking excess returns on the stock market (the value-weighted excess return from Kenneth French). The strong correlation between the two series (0.68) is immediately apparent. Both high- and low-frequency movements in risk aversion are associated with changes in expeted returns. For example, we see the broad pattern that risk aversion is high from the late 1970's to the early 1990's, and future returns are also high in those periods. At higher frequencies, we can see, for example, that the market decline and recession in 2000 and 2001 is associated with both an increase in risk aversion and an increase in future returns.

To see how the predictive power varies with horizon and compares with other variables, figure 1 plots R^2s from regressions of future stock returns on $\hat{\alpha}_t$, cay_t, the price-dividend ratio (P/D), and the excess consumption ratio from Campbell and Cochrane (1999). Each line gives the R^2 from a univariate regression. The x-axis gives the horizon for the return in quarters. The n^{th} point is the R^2 from a regression of $\sum_{j=1}^{n} r_{t+j}$ on the predictor at time t. The regressions are all run on quarterly data from 1952 to 2001 (to ensure that we have data for the 40-quarter regression). Each regression uses the same sample for the predictors.

At every horizon, $\hat{\alpha}$ is dominant. At the five-year horizon, the R^2 for estimated risk aversion

the results are robust to different choices for ϕ.

peaks at 50 percent, nearly than twice that of the other variables. At ten years, the R^2 is still 37 percent. Furthermore, in horse-race regressions (reported in the appendix), $\hat{\alpha}$ dominates the strongest of the other variables, *cay*, at all horizons.

An important consideration in long-horizon forecasting regressions is that the residuals are highly persistent. Kiefer and Vogelsang (2005) show that by using Newey–West standard errors with a very long lag window, we can obtain test statistics with better size properties than techniques that use a fixed (and usually short) lag window. I choose a lag window equal to half the sample size and use the critical values reported in Kiefer and Vogelsang (2005). For *cay*, every regression except for those with horizons greater than 30 quarters is significant at the 5 percent level. For $\hat{\alpha}$, the *largest* p-value is 0.0008. The price-dividend ratio is significant at the 5 percent level for forecasts of 14 quarters or longer. In other words, these regressions all imply that we have substantial ability to forecast stock returns in the postwar period, and $\hat{\alpha}$ is the strongest of the predictors. Out-of-sample tests with both asymptotic and bootstrapped critical values give similar results (appendix D.3). Appendix D examines the sensitivity of the results in this section to the various parameters that are calibrated (e.g., the EIS and the persistence of habits). The basic results hold across a broad range of parameter sets.

Figures 1 and 3 together show that the novel forecasting variable implied by the EZ-habit model outperforms by a wide margin the other major return forecasting variables studied in the literature.

4.3 Forecasts from estimates of technology

The method of estimating the level of risk aversion studied above does not rely on any assumptions about the structure of production in the economy, being derived purely from the preference specification. However, in the production model, changes in lifetime utility are closely related to changes in productivity. If we can measure innovations to technology, then risk aversion should follow an AR(1) process in which the innovations are equal to the shocks to the stochastic trend in technology.

The literature on estimating aggregate technology shocks is extensive. I consider two methods here. The first follows Solow (1957) and uses restrictions from a constant-returns production function:

$$a_t = y_t - \gamma k_t - (1 - \gamma)\, l_t \tag{23}$$

a_t measures technology if the economy has a Cobb–Douglas production function, with y denoting log output, k log capital and l log labor supply.[33] I also consider a simpler metric, labor productivity, $lp_t = y_t - l_t$. Labor productivity does not take into account the effects of capital accumulation and simply models technology as the average product of labor. Capital can be difficult to measure, whereas the number of hours supplied in the economy is a fairly concrete quantity (though the quality of those hours is difficult to account for).[34]

To extract the stochastic trend from the two productivity series, I estimate univariate ARMA models for each variable. The Bayesian information criterion implies that TFP growth is best fit with an MA(2), while labor productivity growth should be treated as i.i.d. ε_t^{TFP} is defined as the residual in the MA(2), while ε_t^{LP} is simply equal to labor productivity growth. That is, ε_t^{TFP} and ε_t^{LP} are innovations to the Beveridge–Nelson (1981) trends in productivity.

In a log linearization of the production model, risk aversion follows an AR(1) process of the form,

$$\alpha_t = (1-\theta)\,\bar{\alpha}_t + \theta\alpha_{t-1} + \varepsilon_t^X \tag{24}$$

where ε^X denotes a measure of the innovation to the stochastic trend of technology. We then have two measures of α_t, which I denote $\hat{\alpha}^{TFP}$ and $\hat{\alpha}^{LP}$, using ε_t^{TFP} and ε_t^{LP}, respectively.[35] The two measures turn out to be highly correlated (93 percent).

The bottom panel of figure 3 plots five-year excess returns against $\hat{\alpha}^{TFP}$ and $\hat{\alpha}^{LP}$. The two series are both clearly highly correlated with future excess returns (0.45 and 0.48, respectively). The p-values in regressions of quarterly excess returns on $\hat{\alpha}^{TFP}$ and $\hat{\alpha}^{LP}$ are 0.032 and 0.026 (using Kiefer, Vogelsang, and Bunzel, 2000, t-type-statistics to account for autocorrelation). The relation among the three series is clearest around the turning points. Productivity growth begins slowing down around 1970, driving risk aversion upwards. Forward-looking stock returns reach

[33] I obtain highly similar results using the measure of total factor productivity suggested by Basu, Fernald, and Kimball (2006).

[34] Furthermore, labor productivity determines the tradeoff that households face between consumption and leisure. If the capital stock rises because foreigners want to invest more in the Uinted States, household welfare will increase even if TFP does not. Similarly, a tax increase that reduced desired saving could lower welfare and labor productivity, without affecting TFP. And welfare is the relevant input in estimating $\hat{\alpha}_t$.

[35] Note that $\hat{\alpha}^{TFP}$ includes some forward-looking information because its construction requires the estimation of an MA(2) on the full sample. $\hat{\alpha}^{LP}$ does not suffer from this flaw. In both cases we do have to estimate mean productivity growth, but shifts in the estimted mean simply correspond to shifts in the mean of $\hat{\alpha}_t$; they have no effect on its dynamics. In regressions of returns on $\hat{\alpha}_t$, the constant will thus always absorb shifts in $\bar{\alpha}$, so the estimation of the mean of productivity growth is irrelevant for forecasting returns.

their trough on nearly the same date. Productivity growth rises again starting in the mid-1990's, which is exactly when stock returns begin falling again.

The result that productivity forecasts equity returns is novel to this paper. It points to a direct link from production and the real economy to stock returns. Perhaps more importantly, this result is easily explained by the EZ-habit model, but is not predicted by models of time-varying disaster risk. Figure 3b thus provides evidence in favor of the EZ-habit model over models of time-varying disaster risk.

5 Conclusion

This paper presents a model of time-varying risk aversion. It simultaneously matches the basic behavior of macroeconomic and financial aggregates, generating both a high and volatile equity premium and a realistic degree of return predictability. The EZ-habit model gives a framework in which consumption, output, and investment growth are all realistically volatile in both the short- and long-run, consumption growth is nearly a random walk, and risk premia are high and volatile.

More generally, this paper provides a general framework for modeling time-varying discount rates that can be used with other macro models. While I study a simple RBC model here, the preference specification is highly tractable and allows models with more realistic descriptions of the production side of the economy to also accommodate variation in risk premia. It is thus an advance towards being able to study the interaction of financial markets and the real economy.

An obvious next step is to study the EZ-habit preferences in a richer setting. Dew-Becker (2013) estimates a standard medium-scale DSGE model with sticky prices and wages, but with the added feature that risk aversion varies over time, as here. Complementing the results in this paper on equity pricing, Dew-Becker (2013) shows that the EZ-habit model, when augmented with a model of inflation, can match the behavior of the nominal term structure well, generating a strongly upward-sloping term structure of nominal interest rates and a volatile term premium.

References

Abel, Andrew B., "Asset Prices under Habit Formation and Catching up with the Joneses," *The American Economic Review, Papers and Proceedings*, 1990, *80(2)*, 38–42.

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 2004, *59* (4), 1481–1509.

Barberis, Nicholas, Ming Huang, and Tano Santos, "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 2001, *116(1)*, 1–53.

Basu, Susanto, John G. Fernald, and Miles S. Kimball, "Are Technology Improvements Contractionary?," *American Economic Review*, 2006, *96(5)*, 1418–1448.

Beeler, Jason and John Y. Campbell, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 2012, *1(1)*, 141–182. Working paper.

Beveridge, Stephen and Charles R. Nelson, "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the "Business Cycle"," *Journal of Monetary Economics*, 1981, *7(2)*, 151–174.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher, "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 2001, *91(1)*, 149–166.

Brunnermeier, Markus and Stefan Nagel, "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation," *The American Economic Review*, 2008, *98(3)*, 713–736.

Bucciol, Alessandro and Raffaele Miniaci, "Household Portfolios and Implicit Risk Preference," *Review of Economics and Statistics*, 2011, *93(4)*, 1235–1250.

Callen, Michael, Mohammad Isaqzadeh, James D. Long, and Charles Sprenger, "Violent Trauma and Risk Preference: Aretefactual and Experimental Evidence from Afghanistan," 2013. Forthcoming, Quarterly Journal of Economics.

Calvet, Laurent E. and Paolo Sodini, "Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios," 2010. Working Paper.

——, **John Y. Campbell, and Paolo Sodini**, "Fight Or Flight? Portfolio Rebalancing by Individual Investors," *The Quarterly Journal of Economics*, 2009, *124(1)*, 301–348.

Cameron, Lisa and Manisha Shah, "Risk-Taking Behavior in the Wake of Natural Disasters," 2013. Working paper.

Campanale, Claudio, Rui Castro, and Gian Luca Clementi, "Asset Pricing in a Production Economy with Chew-Dekel Preferences," *Review of Economic Dynamics*, 2010, *13(2)*, 379–402.

Campbell, John Y., "Consumption-Based Asset Pricing," in "Handbook of the Economics of Finance," Vol. 1, Elsevier Science, 2003, pp. 803–887.

___ **and John H. Cochrane**, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 1999, *107* (2), 205–251.

___ **and N. Gregory Mankiw**, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual*, 1989.

___ **and Robert J. Shiller**, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1(3)* (3), 195–228.

Carroll, Christopher D., "Portfolios of the Rich," in Luigi Guiso, Michael Haliassos, and Tullio Japelli, eds., *Household Portfolios*, MIT Press, 2002.

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michael Marechal, "Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals," 2013. Working paper.

Constantinides, George M., "Habit Formation: A Resolution of the Equity Premium Puzzle," *The Journal of Political Economy*, 1990, *98(3)*, 519–543.

Dew-Becker, Ian, "Bond Pricing With a Time-varying Price Of Risk in an Estimated Medium-Scale Bayesian DSGE Model," 2013. Working paper.

Dynan, Karen E., "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 2000, *90(3)*, 391–406.

Epstein, Larry G. and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, *57(4)*, 937–969.

——— **and** ——— , “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis,” *The Journal of Political Economy*, 1991, *99(2)*, 263–286.

Evans, Martin, “Real Rates, Expected Inflation and Inflation Risk Premia,” *Journal of Finance*, 1998, *53(1)*, 187–218.

Gertner, Robert, “Game Shows and Economic Behavior: Risk-Taking on "Card Sharks",” *Quarterly Journal of Economics*, 1993, *108(2)*, 507–521.

Gourio, Francois, “Disaster Risk and Business Cycles,” *American Economic Review*, 2012, *102(6)*, 2734–2766.

Graeve, Ferre De, Maarten Dossche, Marina Emiris, Henri Sneessens, and Raf Wouters, “Risk premiums and macroeconomic dynamics in a heterogeneous agent model,” *Journal of Economic Dynamics and Control*, 2010, *34(9)*, 1680–1699.

Grossman, Sanford J. and Robert J. Shiller, “The Determinants of the Variability of Stock Market Prices,” *American Economic Review, Papers and Proceedings*, 1981, *71(2)*, 222–227.

Gruber, Jonathan, “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution,” 2006. NBER Working Paper 11945.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, “Time-Varying Risk Aversion,” 2011.

Guvenen, Fatih, “A Parsimonious Macroeconomic Model for Asset Pricing,” *Econometrica*, 2009, *77(6)*, 1711–1750.

Hall, Robert E., “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 1988, *96(2)*, 339–357.

Hansen, Lars Peter and Ravi Jagannathan, “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, 1991, *99(2)*, 225–262.

Jermann, Urban J., “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, 1998, *41* (2), 257–275.

Judd, Kenneth L., *Numerical Methods for Economists*, MIT Press, Cambridge, MA, 1999.

Kaltenbrunner, Georg and Lars A. Lochstoer, "Long-Run Risk Through Consumption Smoothing," *The Review of Financial Studies*, Forthcoming.

Kiefer, Nicholas M. and Timothy J. Vogelsang, "A New Asymptotic Theory for Heteroskedasticity-Autocorrelation Robust Tests," *Econometric Theory*, 2005, *21(6)*, 1130–1164.

___ , ___ , and Helle Bunzel, "Simple Robust Testing of Regression Hypotheses," *Econometrica*, 2000, *68(3)*, 695–714.

Kreps, David M. and Evan L. Porteus, "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 1978, *46(1)*, 185–200.

LeRoy, Stephen F. and Richard D. Porter, "The Present-Value Relation: Tests Based on Implied Variance Bounds," *Econometrica*, 1981, *49(3)*, 555–574.

Lettau, Martin and Harald Uhlig, "Can Habit Formation be Reconciled with Business Cycle Facts?," *Review of Economic Dynamics*, 2000, *3(1)*, 79–99.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A puzzle," *Journal of Monetary Economics*, 1985, *15(2)*, 145–161.

Melino, Angelo and Alan X. Yang, "State-Dependent Preferences can Explain the Equity Premium Puzzle," *Review of Economic Dynamics*, 2003, *6(4)*, 806–830.

Paravisini, Daniel, Veronica Rappoport, and Enrichetta Ravina, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios," 2013. Working paper.

Post, Thierry, Martijn J. van den Assem, Guido Baltussen, and Richard H. Thaler, "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show," *American Economic Review*, 2008, *98(1)*, 38–71.

Rudebusch, Glenn D. and Eric T. Swanson, "Examining the Bond Premium Puzzle with a DSGE Model," *Journal of Monetary Economics*, 2008, *55(Supplement 1)*, S111–S126.

Shiller, Robert J., “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 1981, *71(3)*, 421–436.

Solow, Robert M., “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics*, 1957, *39(3)*, 312–320.

Tallarini, Thomas D., “Risk-Sensitive Real Business Cycles,” *Journal of Monetary Economics*, 2000, *45(3)*, 507–532.

Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen, “Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam,” *American Economic Review*, 2010, *100(1)*, 557–571.

Vissing-Jorgensen, Anette and Orazio P. Attanasio, “Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion,” *The American Economic Review*, 2003, *93(2)*, 383–391.

Vissing-Jorgenson, Anette, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, *110(4)*, 825–853.

Wachter, Jessica A., “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?,” 2010. Working Paper.

Weil, Philippe, “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, 1989, *24(3)*, 401–421.

Yang, Wei, “Intertemporal Substitution and Equity Premium: A Perspective With Habit in Epstein-Zin Preferences,” 2008. Working Paper.

A The certainty equivalent

This section looks at the relationship between the certainty equivalents using G^{habit} and G^{TV}. I first show that the two certainty equivalents are equal up to a second order approximation around the non-stochastic version of the model. Next, I show that in the continuous-time limit, the preferences associated with the two certainty equivalents are identical.

A.1 Second-order approximation

This section approximates the certainty equivalent $G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)$ where $V_{t+1} = V_t \times \left(1+\sigma\varepsilon_{t+1}\right)$ around the point $\sigma = 0$. We assume that $E_t\varepsilon_{t+1} = 0$ and $E_t\varepsilon_{t+1}^2 = 1$.

Now consider the derivative of $G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)$ with respect to σ,

$$\frac{d}{d\sigma}G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right) = \frac{\frac{d}{d\sigma}E_t\left(G\left(V_{t+1}\right)\right)}{G'\left(G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right)} \tag{A.1}$$

We have

$$\frac{d}{d\sigma}E_t\left(G\left(V_{t+1}\right)\right) = \int G'\left(V_t\left(1+\sigma\varepsilon_{t+1}\right)\right)\varepsilon_{t+1}V_t dF\left(\varepsilon_{t+1}\right) \tag{A.2}$$

where F is the cdf of ε_{t+1}. Evaluated at $\sigma = 0$, $\frac{d}{d\sigma}E_t\left(G\left(V_{t+1}\right)\right) = 0$, and therefore $\frac{d}{d\sigma}G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right) = 0$. So all certainty equivalents taking this form are identical up to the first order in approximations around σ.

Next, consider the second derivative,

$$\frac{d^2}{d\sigma^2}G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right) = \frac{\left[G'\left(G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right)\right]\frac{d^2}{d\sigma^2}E_t\left(G\left(V_{t+1}\right)\right) - \left[\frac{d}{d\sigma}E_t\left(G\left(V_{t+1}\right)\right)\right]\left[\frac{d}{d\sigma}G'\left(G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right)\right]}{\left[G'\left(G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right)\right]^2} \tag{A.3}$$

Since $\frac{d}{d\sigma}E_t\left(G\left(V_{t+1}\right)\right)$ is equal to zero at $\sigma = 0$, we can ignore the second term in the numerator. The second derivative of the expectation is

$$\frac{d^2}{d\sigma^2}E_t\left(G\left(V_{t+1}\right)\right) = \int G''\left(V_t\left(1+\sigma\varepsilon_{t+1}\right)\right)\varepsilon_{t+1}^2V_t^2 dF\left(\varepsilon_{t+1}\right) \tag{A.4}$$

At $\sigma = 0$, $\left.\frac{d^2}{d\sigma^2}E_t\left(G\left(V_{t+1}\right)\right)\right|_{\sigma=0} = G''\left(V_t\right)V_t^2$. We also have $\left.G'\left(G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right)\right|_{\sigma=0} = G'\left(V_t\right)$, and hence $\left.\frac{d^2}{d\sigma^2}G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)\right|_{\sigma=0} = \frac{G''\left(V_t\right)V_t^2}{G'\left(V_t\right)}$. So any two choices of G, say G_1 and G_2 are equivalent up to the second order if $\frac{G_1''\left(V_t\right)}{G_1'\left(V_t\right)} = \frac{G_2''\left(V_t\right)}{G_2'\left(V_t\right)}$ for any V_t. That relationship holds for G^{habit} and G^{TV}.

A.2 Continuous time

Duffie and Epstein (1992) show how to extend Epstein–Zin preferences to continuous time. They derive a utility function following the process

$$dV_t = \mu_t + \sigma_t dB_t \tag{A.5}$$

$$= \left(-f\left(c_t, V_t\right) - \frac{1}{2}A\left(V_t\right)\sigma_t\sigma_t'\right)dt + \sigma_t dB_t \tag{A.6}$$

for a Wiener process dB_t.

As in the main text, suppose the household's certainty equivalent under discrete-time Epstein–Zin preferences is $G^{-1}\left(E_t\left(G\left(V_{t+1}\right)\right)\right)$. Duffie and Epstein (1992) show that the analogous choice of A, obtained as a limiting case as the length of time periods approaches zero, is $A\left(V_t\right) = \frac{G''\left(V_t\right)}{G'\left(V_t\right)}$. In the case where $G^{power}\left(V_t\right) = V_t^{1-\alpha}$, we have

$$A^{power}\left(V_t\right) = \frac{G^{power\prime\prime}\left(V_t\right)}{G^{power\prime}\left(V_t\right)} = \frac{-\alpha}{V_t} \tag{A.7}$$

and for $G^{habit} = (V_t - H_t)^{1-\alpha}$

$$A^{habit}(V_t) = \frac{-\alpha}{V_t - H_t} \tag{A.8}$$

For $G^{TV} = V_t^{1-\alpha_t}$,

$$A^{TV}(V_t) = \frac{-\alpha_t}{V_t} \tag{A.9}$$

So G^{TV} and G^{habit} are identical if $\alpha_t = \alpha \frac{V_t}{V_t - H_t}$, which is what is used in the text.

For all three choices of the certainty equivalent G, we can use the standard choice for f, $f(c_t, V_t) = \frac{\beta}{1-\rho} \frac{c_t^{1-\rho} - V_t^{1-\rho}}{V_t^{-\rho}}$. ρ then determines the elasticity of intertemporal substitution, while A determines risk aversion.

B Derivation of the SDF

We can obtain the stochastic discount factor (SDF) by calculating the intertemporal marginal rate of substitution. We calculate two derivatives. First,

$$\frac{\partial V_t}{\partial C_t} = V_t^{\rho} (1 - \exp(-\beta)) C_t^{-\rho} \tag{B.1}$$

Next, we differentiate V_t with respect to $C_{t+1}(w)$, where w denotes one state of the world, and π_w is the probability of that state,

$$\frac{\partial V_t}{\partial C_{t+1}(w)} = \pi_w V_t^{\rho} \exp(-\beta) R_t^{-\rho} G_t^{(-1)\prime} (E_t [G_t (V_{t+1}(w))]) G_t' (V_{t+1}(w)) V_{t+1}^{\rho}(w) (1-\beta) C_{t+1}^{-\rho}(w) \tag{B.2}$$

where G_t' is the derivative of G_t and $G_t^{(-1)\prime}$ the derivative of G_t^{-1}. $R_t \equiv G^{-1}(E_t G_t (V_{t+1}))$. The subscripts on G_t refer to the fact that G_t depends on the potentially-time-varying parameter α_t. The assumption that α_t is exogenous to the household is necessary for this formula for the derivative to be correct (in the same way that external habits lead to a more tractable formula for the SDF than do internal habits).

The SDF can be derived from a consumer's first order conditions for optimization as $M_{t+1}(w) = \frac{1}{\pi_w} \frac{\partial V_t / \partial C_{t+1}(w)}{\partial V_t / \partial C_t}$. We then have

$$M_{t+1}(w) = \exp(-\beta) \frac{G_t'(V_{t+1}(w))}{G_t'(R_t)} \frac{V_{t+1}^{\rho}(w)}{R_t^{\rho}} \frac{C_{t+1}^{-\rho}(w)}{C_t^{-\rho}} \tag{B.3}$$

where the last line follows from the fact that $G_t^{(-1)\prime}(x) = 1/G_t'(x)$.

In the case of $G_t(V) = V^{1-\alpha_t}$, the SDF becomes

$$M_{t+1} = \exp(-\beta) \frac{V_{t+1}^{\rho-\alpha_t}(w)}{R_t^{\rho-\alpha_t}} \frac{C_{t+1}^{-\rho}(w)}{C_t^{-\rho}} \tag{B.4}$$

B.1 Substituting in an asset return

Consider an asset that pays C_t as its dividend. We guess that its cum-dividend price is $W_t = V_t^{1-\rho} C_t^{\rho} (1 - \exp(-\beta))^{-1}$. This guess can be confirmed by simply inserting it into the household's Euler equation.

The return on the consumption claim is

$$R_{w,t+1} = \frac{W_{t+1}}{W_t - C_t} = \frac{V_{t+1}^{1-\rho}}{\exp(-\beta) R_t (V_{t+1})^{1-\rho}} \left(\frac{C_{t+1}}{C_t} \right)^{\rho} \tag{B.5}$$

Which yields

$$\frac{V_{t+1}^{\rho-\alpha_t}(w)}{R_t^{\rho-\alpha_t}} = (R_{w,t+1} \exp(-\beta))^{\frac{\rho-\alpha_t}{1-\rho}} \left(\frac{C_{t+1}}{C_t} \right)^{-\rho \frac{\rho-\alpha_t}{1-\rho}} \tag{B.6}$$

We can then insert this into the SDF to yield

$$M_{t+1} = \exp(-\beta)^{\frac{1-\alpha_t}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\rho\frac{1-\alpha_t}{1-\rho}} R_{w,t+1}^{\frac{\rho-\alpha_t}{1-\rho}} \tag{B.7}$$

C Details of return forecasting

C.1 The method from Lettau and Ludvigson (2001)

If consumption and wealth are cointegrated, then we have the relationship

$$c_t = \zeta w_t + \xi_t \tag{C.1}$$

where ζ is a parameter, and ξ_t is a mean-zero, stationary, and not necessarily i.i.d. error term. If we observed wealth, ζ and ξ_t could be directly estimated. We do not observe wealth, though, especially the human component. Lettau and Ludvigson (2001) therefore use the approximation

$$w_t = \omega a_t + (1-\omega) hu_t \tag{C.2}$$

where a_t is asset wealth and hu_t human wealth. This equation simply says that log aggregate wealth is equal to the sum of log asset and human wealth. Since the level of aggregate wealth is equal to the sum of the levels of asset and human wealth, the approximation is valid as long as the shares of asset and human wealth in aggregate wealth are stationary not not too variable. The fact that labor's share of income has been stationary in the post-war US data makes this assumption reasonable.

Finally, we assume that labor income, y_t, can be viewed as the dividend from human wealth and that the dividend/price ratio for human wealth is stationary. That is,

$$y_t = g + hu_t - \mu_t \tag{C.3}$$

where g is a parameter and μ_t is a mean-zero stationary $b_{z,1}$ term. This implies that

$$w_t = \omega a_t + (1-\omega) y_t + (1-\omega) g + \mu_t \tag{C.4}$$

$$c_t = \zeta\omega a_t + \zeta(1-\omega) y_t + \zeta(1-\omega) g + \zeta\mu_t + \xi_t \tag{C.5}$$

since $\xi_t + \zeta\mu_t$ is mean-zero and stationary, regardless of any correlation between ξ_t and μ_t, the variables c_t, a_t, and y_t are jointly cointegrated. The parameters ζ, ω, and g can be estimated through standard methods for cointegrated models. As Lettau and Ludvigson point out, the estimation is of these parameters is superconsistent, converging linearly with sample size, so these parameters can be taken as known with certainty in any subsequent analyses (in particular, stock return forecasts).

I follow Lettau and Ludvigson in referring to the cointegrating residual, $\zeta\mu_t + \xi_t = c_t - \zeta\omega a_t - \zeta(1-\omega) y_t - \zeta(1-\omega) g$ as cay_t, and I refer to $\omega a_t + (1-\omega) \times y_t$ as ay_t. ay_t is an estimate of total wealth derived from data on consumption, asset wealth, and labor income, taking advantage of an assumed cointegrating relationship between the three variables. I estimate the parameters using standard maximum likelihood methods.

C.2 Sensitivity analysis for return forecasting

The results in section 4.2 depend on choices for two parameters – the EIS and the persistence of risk aversion. Tables A3 and A4 report the ratio of the R^2 for excess value to *cay* for 1, 5, 10, and 20-quarter returns across a variety of choices for the EIS and the persistence of risk aversion.

Table A2 varies the EIS between 0.75 and 10. The numbers in bold represent points where *cay* outperforms $\hat{\alpha}$. When the EIS is greater than 1, *cay* only ever outperforms at the 1-quarter horizon, and then only if the EIS is set to 10. With an EIS less than 1, though, *cay* always has an R^2 substantially larger that of $\hat{\alpha}$. Moreover, the sign on $\hat{\alpha}$ in the return regressions flips. Intuitively, this is because in the construction of $\hat{v}$, when the EIS is less than 1, the weight on aggregate wealth is negative. The theory would predict that high risk aversion is associated with low returns, but with the EIS less than 1, $\hat{\alpha}$ and future returns are actually positively correlated.

Table A3 presents R^2 ratios for the same set of regressions, but now varying the persistence of risk aversion. Across a fairly wide range of autocorrelations, $\hat{\alpha}$ outperforms *cay* at most horizons. The best performance is found with an annual autocorrelation of 0.9, which corresponds to $\phi = 0.974$. Even with an autocorrelation as low as 0.65 ($\phi = 0.9$), though, $\hat{\alpha}$ performs nearly as well as *cay*. As with the EIS, the place where *cay* is most likely to outperform is with 1-quarter returns. Table A4 lists R^2s for *cay*, PE, and $\hat{\alpha}$ for pre and post-1980 samples.

C.3 Out-of-sample forecasting regressions

An alternative to the in-sample regressions studied in the main text is out-of-sample tests of forecasting power. I consider the mean squared forecast $b_{z,1}$ (MSFE) based tests from analyzed in Clark and McCracken (2001, 2005) and Clark and West (2007).

Suppose we want to test whether a single variable, x_t, forecasts stock returns, r_t, against the null that r_t is i.i.d. (the methods used here apply to any null model that is nested; i.e. they are appropriate for asking whether x_t has marginal forecasting power when added to some other model). The forecast horizon can be any length. Therefore, denote $r_{t,t+j} \equiv \sum_{\tau=t}^{t+j} r_\tau$.

We compare the residuals from the null model, $e_{1t} \equiv r_{t,t+j} - \hat{\beta}_{0,t}$ (for an estimated constant mean $\hat{\beta}_{0,t}$ using data prior to date t) to the residuals from the alternative model, $e_{2t,t+j} \equiv r_{t,t+j} - \hat{\beta}_{0,t} - \hat{\beta}_{1,t} x_{t+j-1}$ (where $\hat{\beta}_{1t}$ is a constant regression coefficient estimates on the data from $\tau = 0$ to $\tau = t-1$). The samples for the regressions are begun after the first 20 percent of the sample.

The measure of the difference in MSFE is

$$f_{t,t+j} \equiv e_{1t,t+j}^2 - e_{2t,t+j}^2 + \left(e_{1t,t+j} - e_{2t,t+j}\right)^2 \tag{C.6}$$

Under the null, the MSFE for the e_1 model tends to be smaller than the MSFE for the e_2 model because the e_2 model has added noise due to the extraneous predictor. Intuitively, model e_1 correctly imposes the constraint that $\beta_1 = 0$ under the null. The term $\left(e_{1t,t+j} - e_{2t,t+j}\right)^2$ is essentially a correction for this effect.

When the forecast horizon is more than a single observation, $f_{t,t+j}$ is serially correlated. To correct for this, we divide by a consistent estimate of its long-run variance (spectral density at frequency zero). Following Clark and West (2007), I use the Newey–West measure with a lag window of $1.5 \times j$. Denote this measure of the long-run variance as $S\left(f_{t,t+j}\right)$. The long-run variance corrects for the fact that the forecast $b_{z,1}$ from overlapping samples will be serially correlated. Clark and McCracken tabulate the critical values of the statistic $\left((T-j)\sum_{t=1}^{T-j} f_{t,t+j}\right)/S\left(f_{t,t+j}\right)$.

In the main text, $\hat{\alpha}$ is calculated using full-sample information. In particular, we need to calculate the cointegrating relationship between consumption, labor income, and financial wealth. We also need to know the average growth rate of value. For the out-of-sample forecasts, all of those parameters are estimated using only backward-looking information. The only possible source of look-ahead bias here would be data revisions.

The top panel of figure A2 plots the values of the statistics using $\hat{\alpha}_t$ as the predictor against a null of a constant expected equity for horizons from 1 to 20 quarters. We can easily reject the null at the 5 percent level at all horizons and at the 1% level for 2–13 quarter horizons.

C.3.1 Bootstrapping

A major concern with predictive regressions is that asymptotic distribution theory is often a poor guide to small-sample behavior. A simple way to deal with that concern is to use a bootstrap to construct confidence intervals for the test statistics. I construct bootstrap samples in the following way. I select bootstrap samples of stock returns and growth rates of consumption, asset wealth, and labor income. I then construct level series for consumption, wealth, and income, and calculate $\hat{\alpha}$ using purely backward-looking information as above. Finally, I construct the test statistic from above for each bootstrapped sample at each horizon from 1 to 20 quarters. I bootstrap 10,000 samples of data. The top panel of figure A2 plots the 95th and 99th percentiles of the bootstrapped test statistics, and the out-of-sample forecasting power is still significant at the 5 percent level.

C.3.2 $\hat{\alpha}$ versus *cay*

We can also use the out-of-sample test to ask whether estimated risk aversion forecasts stock returns better than *cay*. The null model is now one where stock returns depend on a constant and the lagged value of *cay*, and the encompassing alternative adds the lagged value of $\hat{\alpha}$. The bottom panel of figure A2 plots the test statistics. At every horizon, we can reject the null that $\hat{\alpha}$ does not improve the forecast using *cay* at the 5 percent level, and we can reject the null at the 1 percent level at every horizon longer than 1 quarter.

Figure A2 also plots the statistic for a test of whether *cay* has any marginal predictive power above that of $\hat{\alpha}$. At horizons shorter than 8 quarters, we cannot reject the null that it does not. At longer horizons, though, there is evidence that both variables contain important information for forecasting stock returns.

D Alternative habit formation models

This section reports basic macro and financial moments for four alternative models that involve various types of habit formation. In each case, I combine the preferences with an RBC model with random-walk technology calibrated as in the main text.

The first model is Campbell–Cochrane (1999) preferences. The specific preference specification is

$$\begin{aligned} U\left(C_t\right) &= \frac{\left(C_t - X_t\right)^{1-\rho}}{1-\rho} && \text{(D.1)} \\ s_t &= \log\left(\left(C_t - X_t\right)/C_t\right) && \text{(D.2)} \\ s_t &= \left(1-\phi\right)\bar{s} + \phi s_{t-1} + \exp\left(-\bar{s}\right)\left(\sqrt{1-2\left(s_t-\bar{s}\right)}-1\right)\left(\Delta c_t - E_{t-1}\Delta c_t\right) && \text{(D.3)} \end{aligned}$$

where s_t is the excess-consumption ratio. The SDF is

$$M_{t+1} = \beta\left(\exp\left(s_{t+1}-s_t\right)\frac{C_{t+1}}{C_t}\right)^{-\rho} \tag{D.4}$$

I calibrate the preferences as in Campbell and Cochrane (1999), with $\rho = 2$, $\phi = 0.87^{1/4}$, and $\bar{s} = \log(0.057)$.

The second model follows that of Jermann (1998). In this case, utility is

$$U\left(C_t\right) = \frac{\left(C_t - b\bar{C}_{t-1}\right)^{1-\rho}}{1-\rho} \tag{D.5}$$

where $\bar{C}$ is aggregate consumption. Jermann's model also involves adjustment costs in investment, which I calibrate as he does. The coefficient b is set to 0.82, and $\rho = 5$. This model thus has two key differences with Campbell and Cochrane (1999): the habit is a smaller fraction of total consumption, and there are adjustment costs in investment, which should make consumption growth and asset prices more volatile.

Table 1. Calibration		
Parameter	**Value**	**Target**
γ	0.33	Capital income share
β	0.9975	2% annual real risk-free rate
δ	0.02	8% annual depreciation (BEA data)
μ	0.005	2% annual output growth
φ	0.94	Persistence of price/dividend ratio
ϱ	0.67	A priori (see text)
σ_a	0.0088	Long-run standard deviation of consumption growth
mean(α_t)	14	Mean Sharpe ratio (0.32 annualized)
stdev(α_t)	6.2	Stock return predictability
σ_x	0.012	Variance of output growth
φ_x	0.9	Variance of output growth

Note: Parameters used for the structural models. In table 2, the CRRA model uses with stdev(α)=0; the benchmark EZ-habit model (column 3) sets σ_x=0.

Table 2. Comparison of preference specifications						
	1	2	3	4	5	6
Model:	Data	EZ-habit	Dual-shock	EZ-CRRA	Campbell–Cochrane	Jermann
Real moments:						
1 Long-run SD(dC,dY,dI) (%)	0.88	0.88	0.88	0.88	0.88	0.88
2 StdDev(dY) (%)	0.99	0.59	1.03	0.59	0.59	0.59
3 StdDev(dC) (%)	0.46	0.47	0.56	0.28	0.15	0.38
4 StdDev(dI) (%)	2.65	0.83	2.37	1.11	2.21	1.32
Financial moments:						
5 Mean SR (annualized)	0.32	0.32	0.32	0.22	0.07	0.17
6 Mean Rk (annualized %)	6.78	4.07	4.05	0.90	-0.13	1.14
7 StdDev(Rk) (annualized %)	21.19	10.67	10.65	4.11	4.48	11.66
8 Mean Rf (annualized %)	0.91	2.04	1.94	2.20	4.99	0.24
9 StdDev(Rf) (annualized %)	1.16	0.25	0.26	0.21	0.11	10.20
Predictability of consumption and returns						
10 corr(dC(t),Rf(t-1))	-0.09	0.07	0.08	0.28	0.75	-0.61
11 corr(dC(t->t+20),PD(t-1))	-0.05	-0.04	0.01	0.21	0.52	0.04
12 corr(Rk(t->t+20),PD(t-1))	-0.41	-0.44	-0.35	-0.25	-0.25	-0.17

Note: All models with Epstein–Zin preferences are calibrated as in table 1, while the last two columns are calibrated as in the original papers. All variables are measured using quarterly values unless otherwise specified. dI is investment growth, dY output growth, and dC consumption growth. Rf is the risk-free rate (measured empirically as the nominal 3-month yield minus an inflation forecast), and Rk is the annualized return on a levered dividend claim (with a leverage ratio of 2.4). The long-run SD is the square root of the spectral density at frequency zero multiplied by 2π. SR is the annualized Sharpe ratio. The correlations in the final three rows use 20-quarter forward-looking consumption growth and equity returns. PD is the price/dividend ratio on the dividend claim.

Table 3. Regressions estimating the elasticity of intertemporal substitution

	Model:	Data	EZ-CRRA	EZ-habit	Dual-shock
1	Population, infeasible ($E_t[\Delta c_{t+1}]$)	N/A	1.50	0.64	0.78
2	Population	N/A	1.50	0.56	0.71
3	Small sample	0.14	1.16	0.03	0.35
4	[2.5%, 97.5%]	N/A	[0.03, 1.79]	[-1.98, 1.02]	[-1.32, 1.34]
5	Population, RRA control	N/A	N/A	1.50	1.50
6	Small sample, RRA control	0.18	N/A	0.07	0.94
7	[2.5%, 97.5%]	N/A	N/A	[-3.08, 3.11]	[-1.07, 2.78]

Note: Values reported are the coefficients from regressions of consumption growth or expected consumption growth on the risk-free rate. The dependent variable in row 1 is expected consumption growth (computed numerically in the simulations); all other rows use realized consumption growth. The small-sample regressions are based on 228 quarters of data, and median coefficient estimates are reported; 2.5 and 97.5 percentiles are reported in brackets. The RRA control is actual risk aversion in the simulations and estimated risk aversion (section 4) in the empirical regressions.

Figure 1. Simulated and empirical R^2s

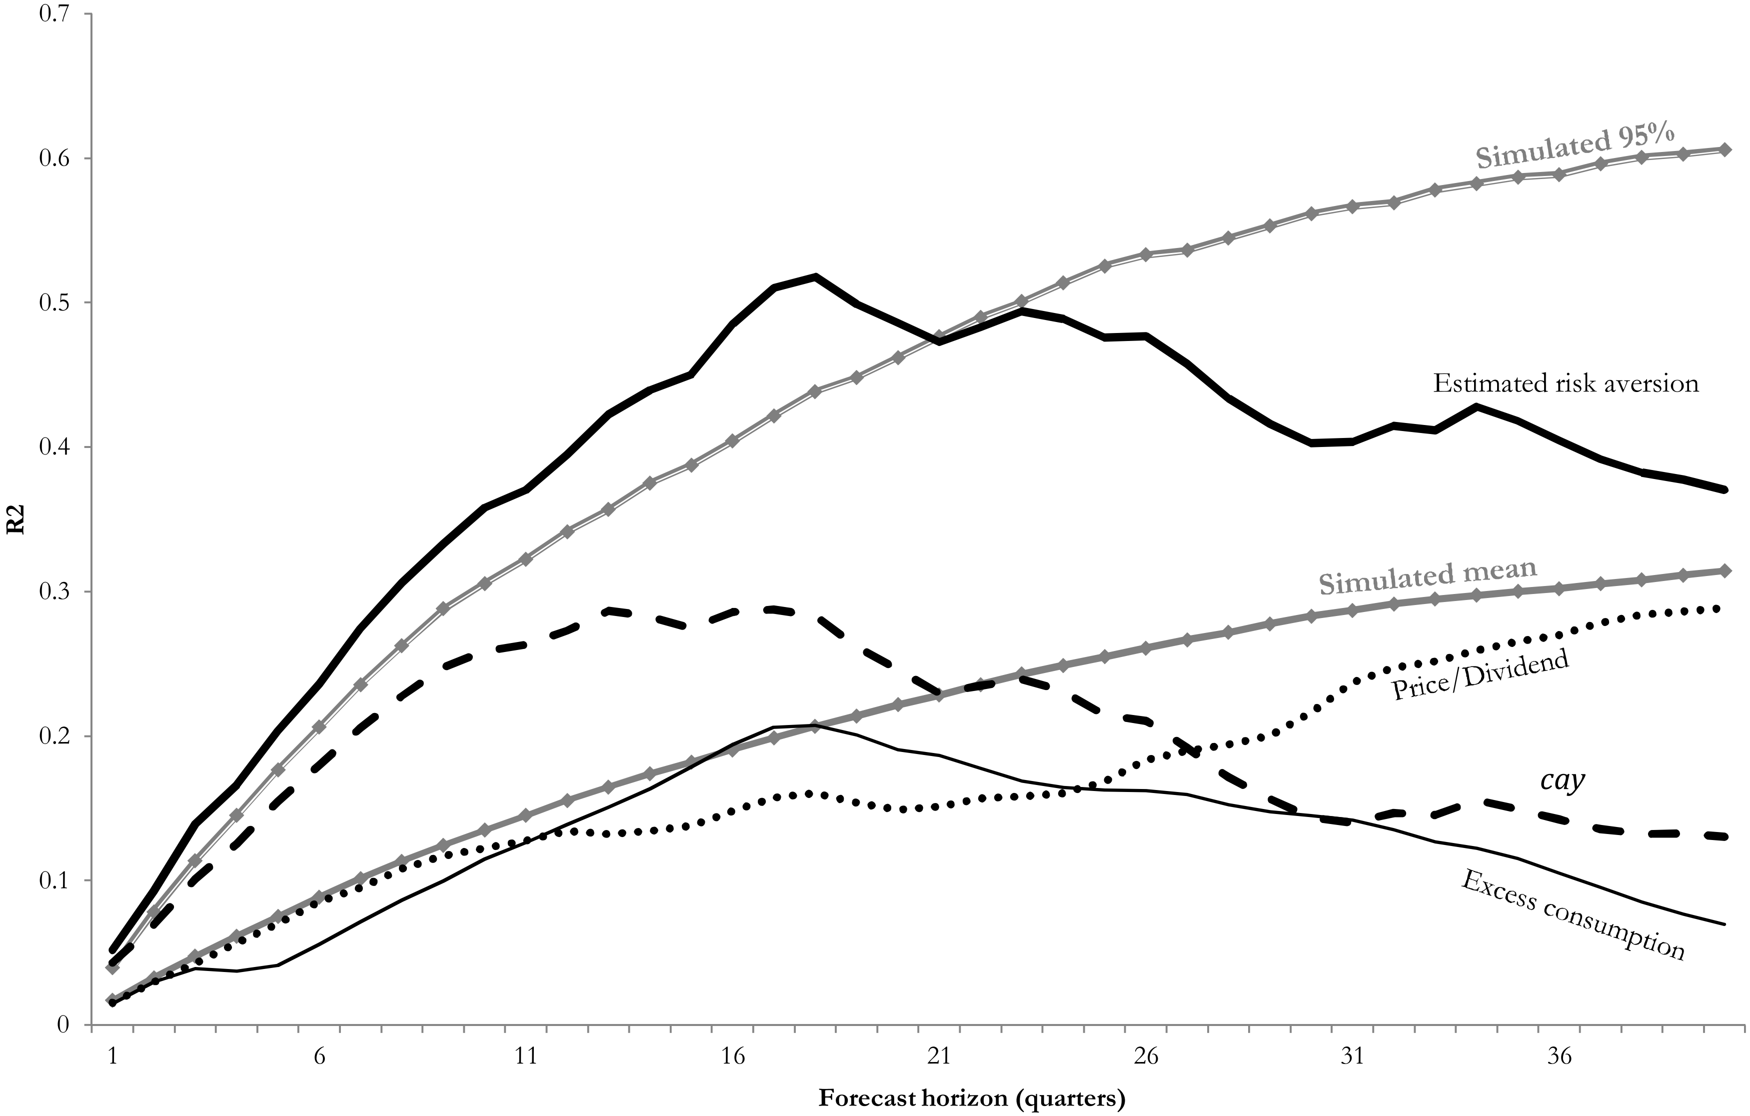

Note: R^2s from univariate regressions of stock returns on various predictors. The forecast horizon is reported in quarters. Data for *cay* is obtained from Sydney Ludvigson's website; Price/dividend data comes from CRSP; the Campbell–Cochrane excess consumption ratio is computed using their parameter values and consumption data from the BEA. The gray lines give the mean and 95th percentile in the simulation of the EZ-habit model.

Figure 2. Impulse response functions

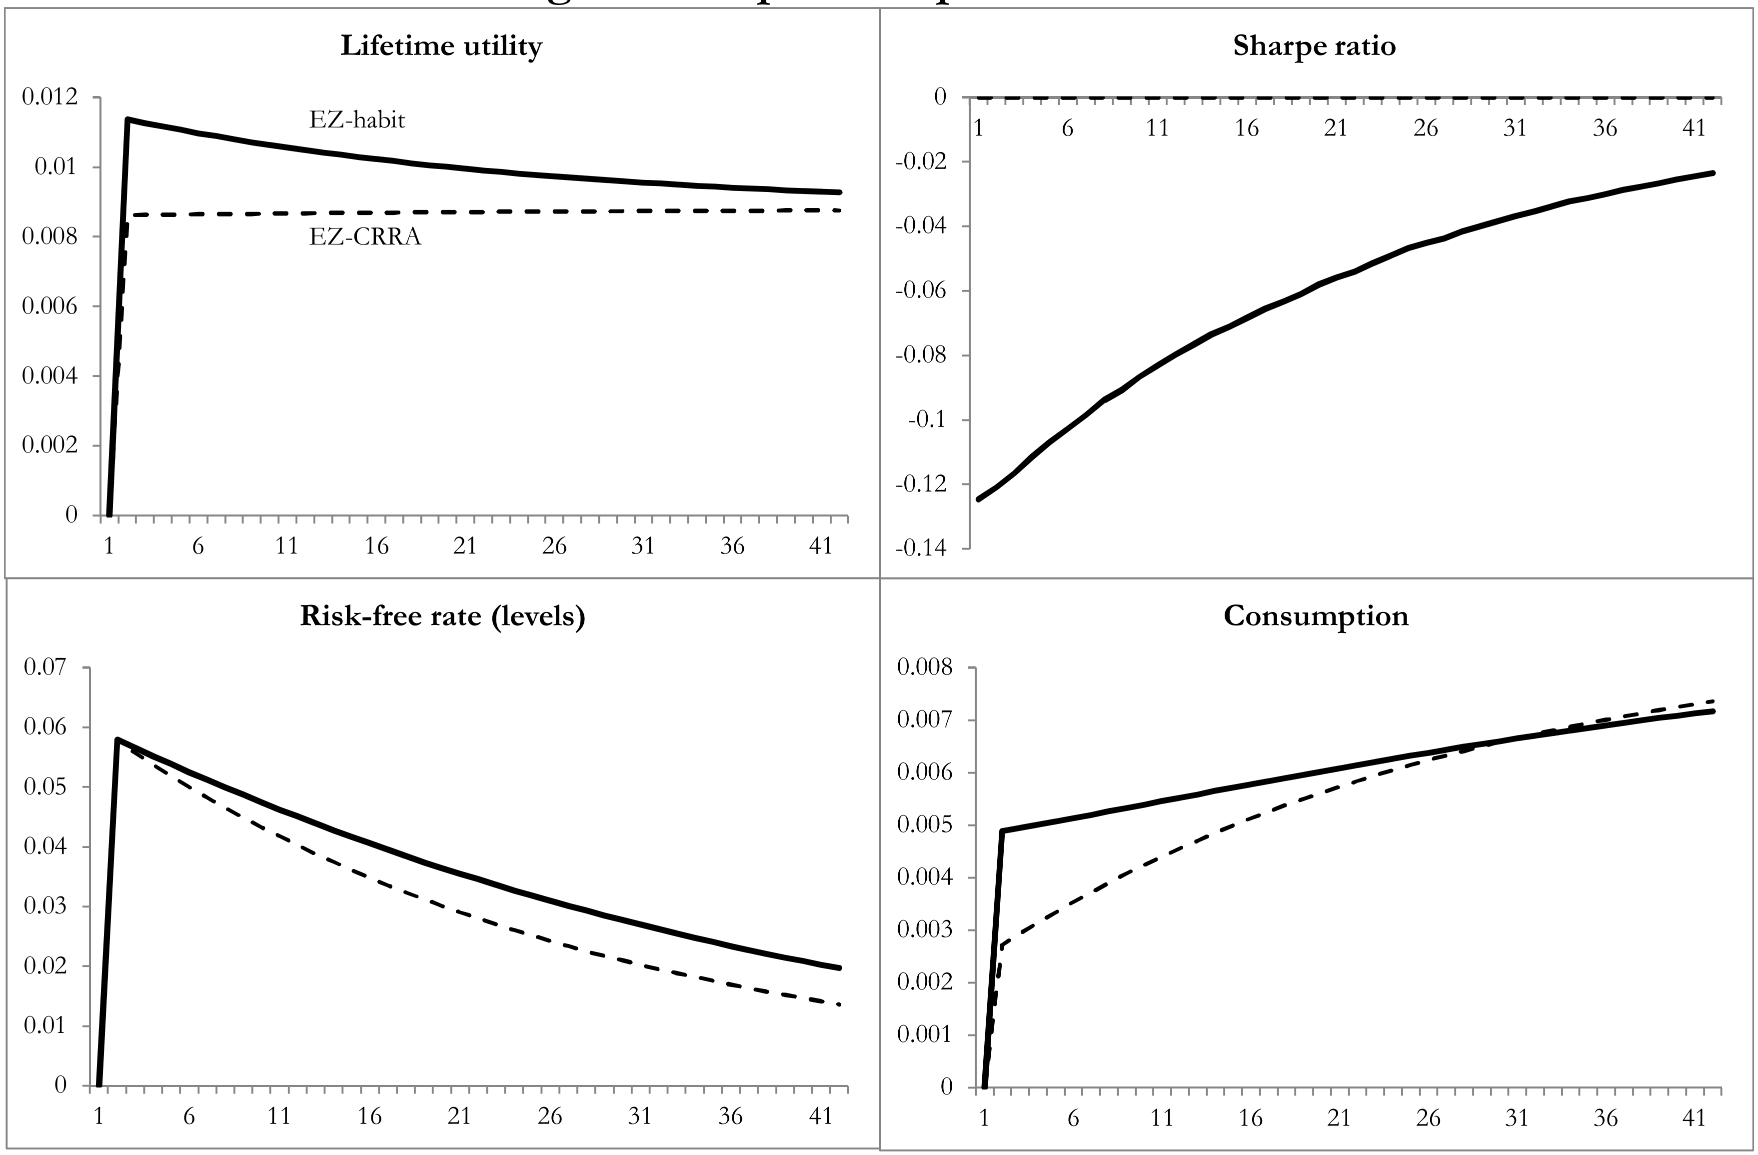

Note: Impulse responses for the EZ-CRRA and EZ-habits models. The shock is a positive unit-standard-deviation increase in technology. The dotted lines are for EZ-CRRA, solid are for EZ-habit. All functions are reported as fractions of the variables' means except for the risk-free rate, for which the response is in annualized percentage points. Value is lifetime utility; the Sharpe ratio is for an asset that pays aggregate consumption as its dividend.

Figure 3a. Estimated risk aversion and 5-year excess stock returns

Note: Excess stock returns are for the CRSP value-weighted index minus the risk-free rate. Returns are forward-looking five-year averages. Risk aversion is estimated from data on aggregate wealth and consumption and is normalized to have zero mean and unit variance.

Figure 3b. Stock returns and estimates of risk aversion from productivity growth

Note: Total factor productivity is the quarterly Solow residual from John Fernald's website. Labor productivity is output per hour in the non-farm private business sector from the BLS. Risk aversion is an AR(1) with innovations equal to the (negative) innovations to the Beveridge–Nelson trend in productivity.

Table A1. Comparison of results from simulations of projection and log-linear model solutions			
Differences:	Mean	Std. dev.	Scaled std. dev.
Capital	-8.72E-04	0.00073	0.020
Cons. Growth	3.16E-08	0.00013	0.028
RRA	1.11E-02	0.75344	0.116
Sharpe Ratio	-2.15E-03	0.00892	0.123

Note: Comparison of the projection and log-linear solutions. The two simulations use the same shocks but different policy functions. The first column is the mean difference between the simulations, the second column the standard deviation, and the third column the standard deviation of the difference scaled by the standard of the variable in the projection solution. RRA is relative risk aversion.

Table A2: Relative R2s for varying EIS

Span	EIS=0.1	0.25	0.75	1.25	1.5	2	10
1 quarter	**0.39**	**0.26**	**0.29**	1.07	*1.10*	1.08	**0.96**
5 quarters	**0.44**	**0.28**	**0.29**	1.16	*1.21*	1.21	1.09
10 quarters	**0.61**	**0.42**	**0.31**	1.41	*1.49*	1.49	1.37
20 quarters	1.28	1.02	**0.21**	1.92	*2.08*	2.15	2.09

Note: This table lists the ratio of the R2 for a univariate regression of long-horizon returns on estimated risk aversion to the R2 for *cay*. Values less than 1 are in bold. The span in quartes is listed in the left hand column. The top row gives the EIS. The EIS is used to calculate household value and risk aversion.

Table A3. Relative R2s for varying persistence of risk aversion

Span	Autocorr.=0.95	0.9	0.85	0.8	0.75	0.7
1 quarter	**0.86**	1.27	*1.10*	**0.90**	**0.75**	**0.64**
5 quarters	1.03	1.44	*1.21*	**0.97**	**0.78**	**0.64**
10 quarters	1.17	1.73	*1.49*	1.20	**0.99**	**0.82**
20 quarters	1.24	2.24	*2.08*	1.76	1.48	1.27

Note: This table lists the ratio of the R2 for a univariate regression of long-horizon returns on estimated risk aversion to the R2 for *cay*. Values less than 1 are in bold. The span in quarters is listed in the left hand column. The top row gives the annual autocorrelation of risk aversion.

Table A4. R2s from pre and post-1980 univariate return forecasting regressions

pre-1980	Estim. RRA	*cay*	P/D	post-1980	Estim. RRA	*cay*	P/D
1q	0.10	**0.10**	0.03	1q	**0.03**	0.03	0.03
5q	**0.28**	0.22	0.25	5q	**0.18**	0.15	0.08
10q	0.27	0.16	**0.27**	10q	**0.48**	0.36	0.19
20q	0.38	0.13	**0.39**	20q	**0.56**	0.33	0.29

Note: R2s from univariate regressions of long-horizon stock returns on estimated risk aversion, *cay*, and the price/dividend ratio. The highest value for each horizon and sample is listed in bold.

Figure A1. Log_{10} Euler equation error densities

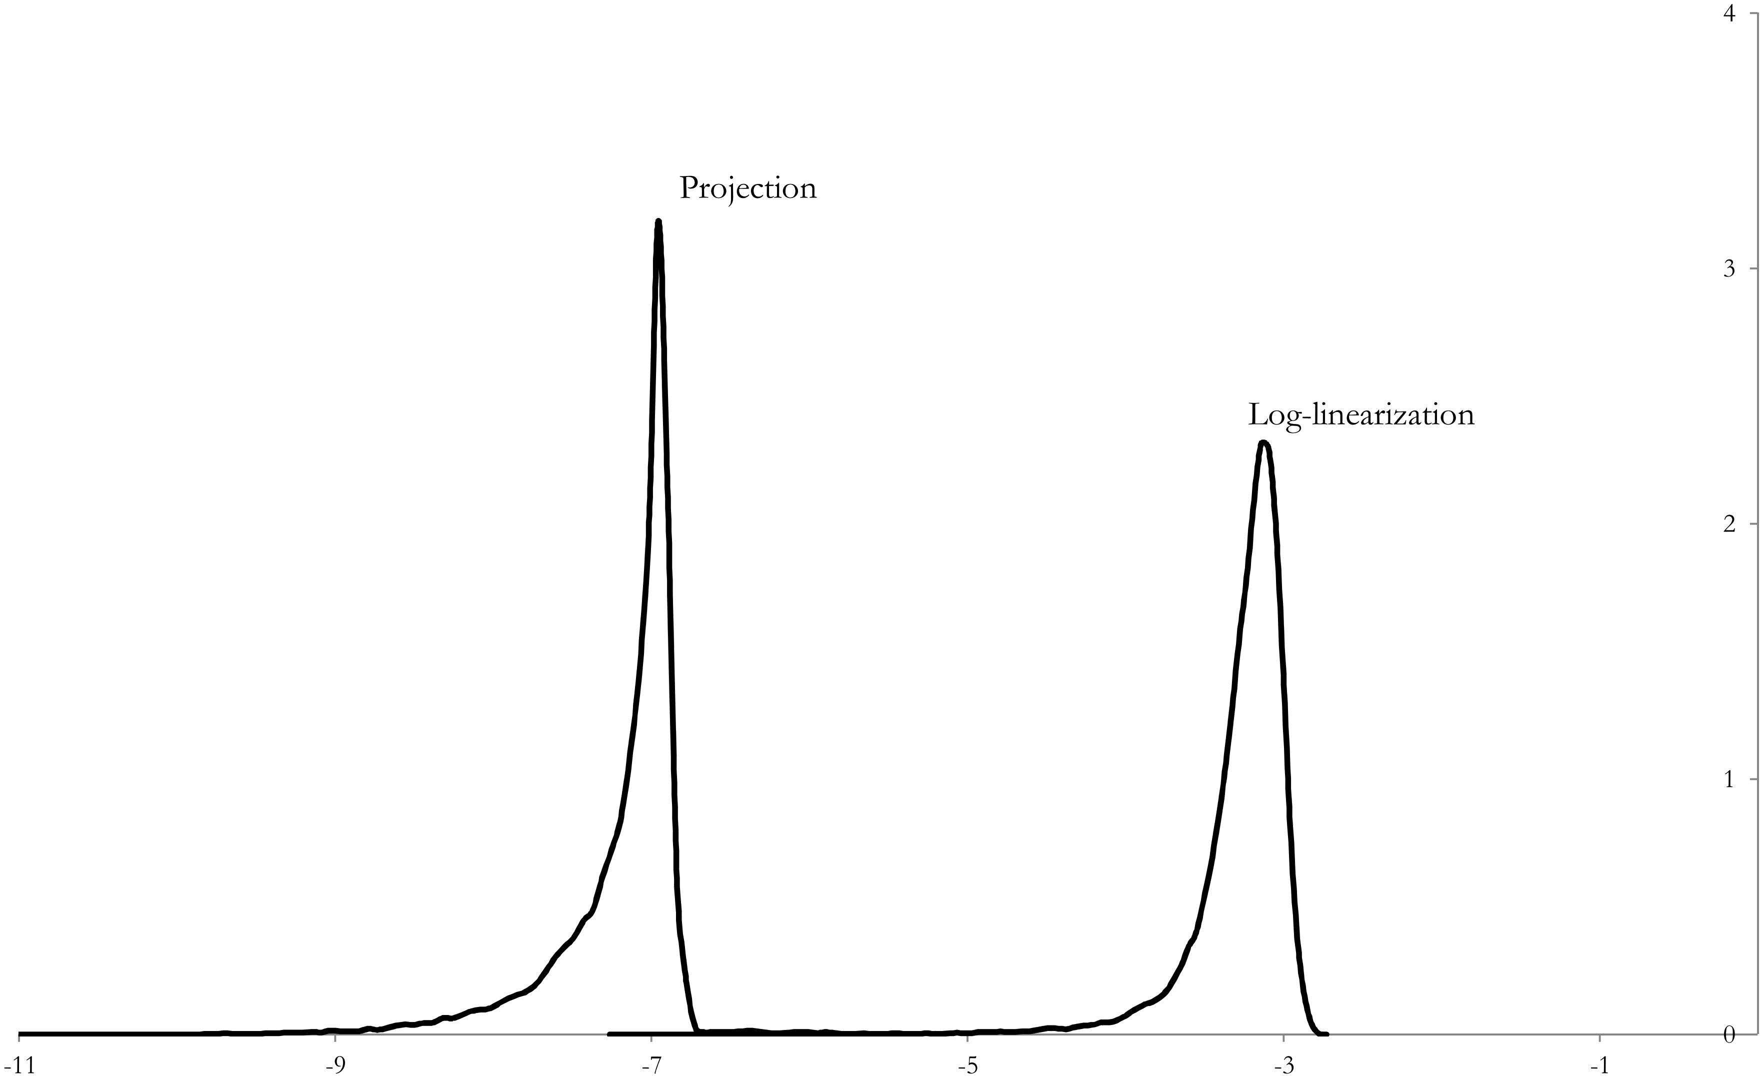

Note: Densities of Euler equation errors under the two solution methods. The log errors are defined as $\log_{10}(|E[M_{t+1}R_{k,t+1}]-1|)$. Densities are estimated using a kernel smoother on simulated data. In both cases, the model used is the benchmark single-shock model with EZ-habit preferences and constant labor supply.

Figure A2. Out-of-sample test statistics

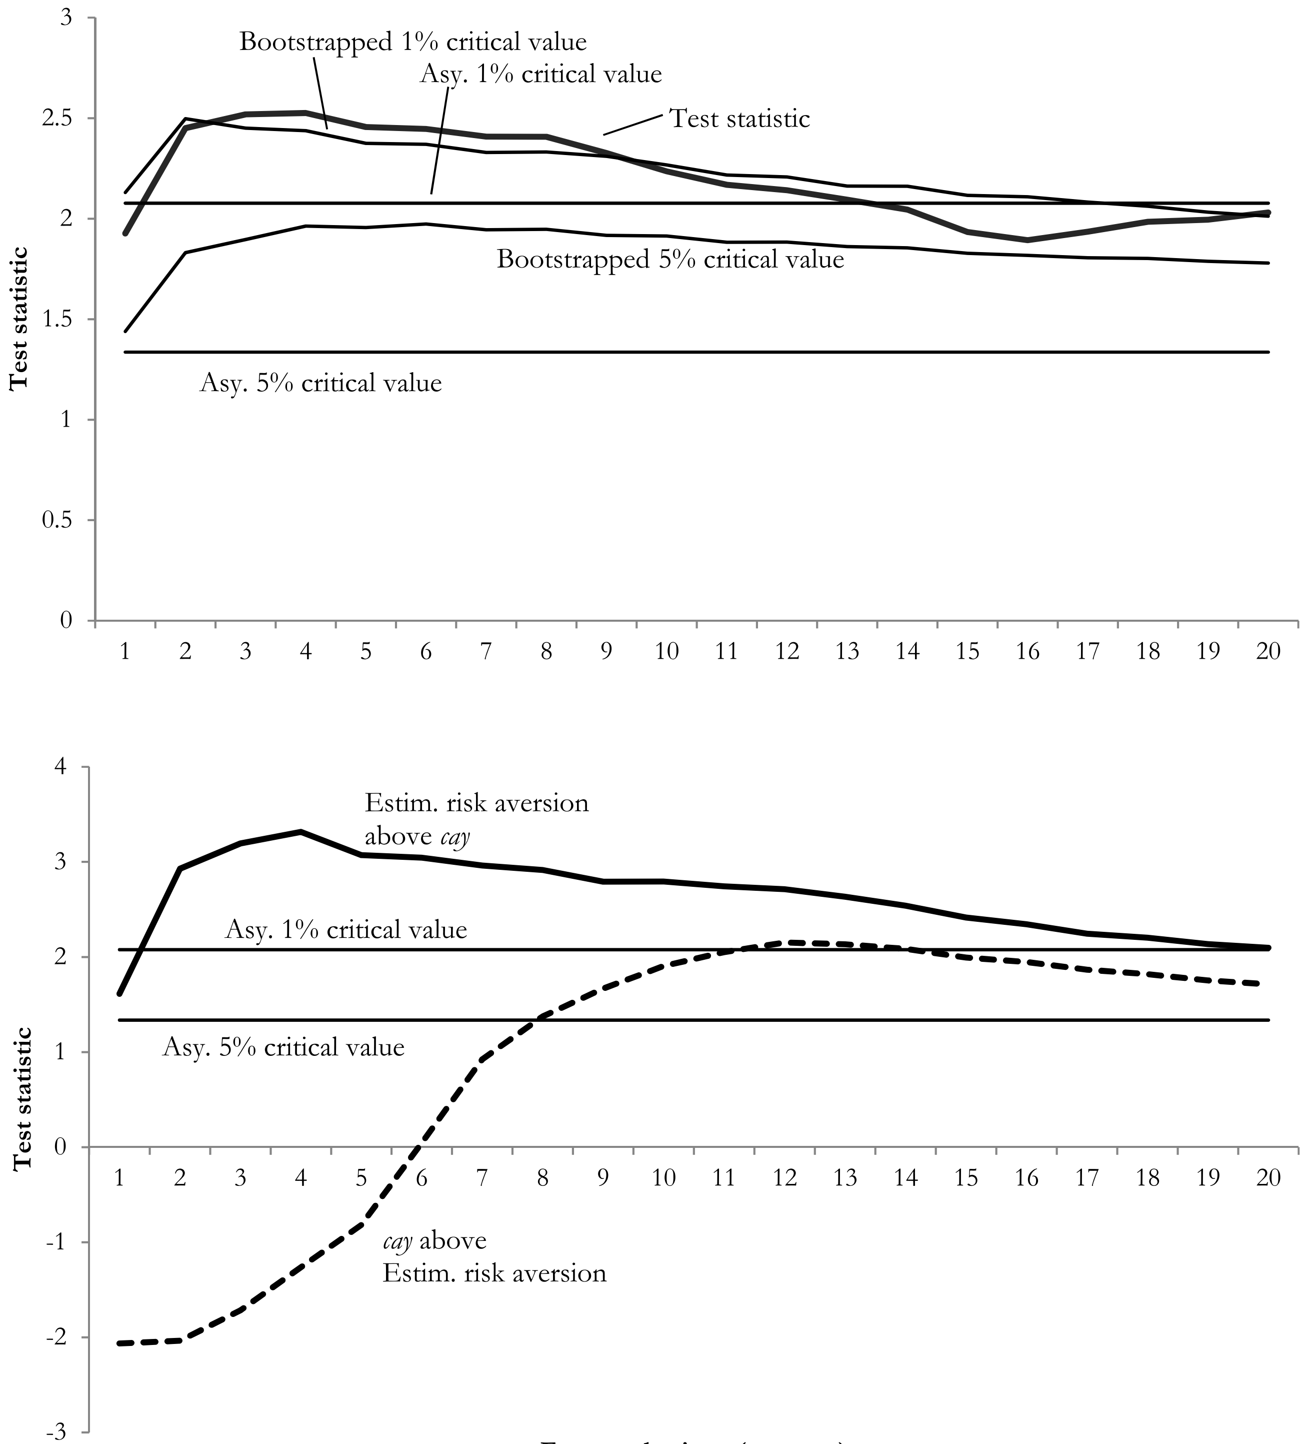

Note: Out-of-sample test statistics from Clark and McCracken (2001, 2005) based on the reduction in out-of-sample RMSE. Estimated risk aversion depends on the cointegrating model used to estimate *cay*. The top panel tests whether estimated risk aversion has marginal forecasting power against a null of a constant-mean model for returns. The cointegrating vector is reestimated in each period using only backward-looking information. The bottom panel tests adding estimated risk aversion to a null model including a constant and cay and vice versa.