The Wealth-Consumption Ratio

Hanno Lustig Stijn Van Nieuwerburgh Adrien Verdelhan*

Abstract

To measure the wealth-consumption ratio, we estimate an exponentially affine model of the stochastic discount factor on bond yields and stock returns. We use that discount factor to compute the no-arbitrage price of a claim to aggregate US consumption. Our estimates indicate that total wealth is much safer than stock market wealth. The consumption risk premium is only 2.2 percent, substantially below the equity risk premium of 6.9 percent. As a result, our estimate of the wealth-consumption ratio is much higher than the price-dividend ratio on stocks throughout the post-war period. The high wealth-consumption ratio implies that the average US household has a lot of wealth, most of it human wealth. A variance decomposition of the wealth-consumption ratio shows less return predictability than for stocks, and some of the return predictability is for future interest rates not future excess returns. We conclude that the properties of the average US household's portfolio are more similar to those of a long-maturity bond than those of stocks. The differences that we find between the risk-return characteristics of equity and total wealth suggest that equity is a special asset class.

*Lustig: Department of Economics, University of California at Los Angeles, Box 951477, Los Angeles, CA 90095; hlustig@econ.ucla.edu; Tel: (310) 825-8018; http://www.econ.ucla.edu/people/faculty/Lustig.html. Van Nieuwerburgh: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; svnieuwe@stern.nyu.edu; Tel: (212) 998-0673; http://www.stern.nyu.edu/~svnieuwe. Verdelhan: Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215; av@bu.edu; Tel: (617) 353-6324; http://people.bu.edu/av. This paper circulated before as "The Wealth-Consumption Ratio: A Litmus Test for Consumption-Based Asset Pricing Models." The authors would like to thank Dave Backus, Geert Bekaert, John Campbell, John Cochrane, Ricardo Colacito, Pierre Collin-Dufresne, Bob Dittmar, Greg Duffee, Darrell Duffie, Lars Peter Hansen, John Heaton, Dana Kiku, Ralph Koijen, Martin Lettau, Francis Longstaff, Sydney Ludvigson, Thomas Sargent, Kenneth Singleton, Stanley Zin, and participants of the NYU macro lunch, seminars at Stanford GSB, NYU finance, BU, the University of Tokyo, LSE, the Bank of England, FGV, MIT Sloan, Purdue, LBS, Baruch, Kellogg, Chicago GSB, and conference participants at the SED in Prague, the CEPR meeting in Gerzensee, the EFA meeting in Ljubljana, the AFA and AEA meetings in New Orleans, and the NBER Asset Pricing meeting in Cambridge for comments. This work is supported by the National Science Foundation under Grant No 0550910.

Stock returns have played a central role in the development of modern asset pricing theory. Yet, in the US, stock market wealth is only a small fraction of total household wealth. Real estate, non-corporate businesses, other financial assets, durable consumption goods, and human wealth constitute the bulk of total household wealth. We measure total wealth and its price-dividend ratio, the wealth-consumption ratio, by computing the no-arbitrage price of a claim to the aggregate consumption stream. To value this claim, we estimate from stock returns and bond yields the prices of aggregate risk that US households face.

We find that the average household's wealth portfolio is more like a long-maturity real bond than like equity, for two reasons. First, the total wealth portfolio earns a low risk premium of around 2.2 percent per year, compared to a much higher equity risk premium of 6.9 percent. As a result, the wealth-consumption ratio is much higher, 87 on average, than the price-dividend ratio on equity, 27 on average. Second, the wealth-consumption ratio is less volatile than the price-dividend ratio: its standard deviation is 17 percent versus 27 percent. The return on total wealth has a volatility that is 9.8 percent per year, compared to 16.7 percent for equity returns. Our estimation produces a variance decomposition of the wealth-consumption ratio in closed form, the no-arbitrage analog to the Campbell and Shiller (1988) decomposition of the price-dividend ratio for stocks. The lower variability in the wealth-consumption ratio indicates less variation in expected future total wealth returns. Hence, there is less predictability in total wealth returns than in equity returns. We find that most of the variation in future expected total wealth returns is variation in future expected risk-free rates, and not variation in future expected excess returns. For the average investor, most of the time variation in his investment opportunity set comes from interest rates, not from risk premia. Finally, we show that the differences between the total wealth portfolio and equity cannot be eliminated simply by thinking of equity as a leveraged claim to aggregate consumption.

The properties of the average household's total portfolio are crucial for the evaluation of dynamic asset pricing theories, business cycle models, and the welfare costs of economic fluctuations. First, Roll (1977) points out that the total wealth return is the right pricing factor in the Capital Asset Pricing Model, while (Campbell 1993) shows that current and future total wealth returns substitute for consumption growth as pricing factors in the Inter-temporal CAPM. However, applied work commonly tests dynamic asset pricing models (DAPMs) by using the stock market return as a proxy for the total wealth return. Second, in the real business cycle literature, the canonical model with log utility implies a constant wealth-consumption ratio: changes in real interest rates are exactly offset by changes in real consumption growth. We document substantial variation in the wealth-consumption ratio we estimate: the changes in predicted consumption growth are too small to offset the effect of changes in interest rates. Third, Alvarez and Jermann (2004) show the mapping between the level of the wealth-consumption ratio and the marginal welfare cost of consumption fluctuations. The per unit benefit of a marginal reduction in aggregate consumption fluctuations is equal to the ratio of the prices of two long-lived securities: one representing a claim to stabilized consumption, the other a claim to actual consumption. Similarly, Martin (2008) derives a simple, analytical relation between the wealth-consumption ratio and the total cost of business cycles, which measures the benefits of completely eliminating aggregate consumption growth risk, in an environment with i.i.d. aggregate consumption growth and power utility. Thus, estimating the price of or the return on a claim to aggregate consumption is a critical challenge for financial and macro-economists. That is the goal of our paper.

In the absence of a clear candidate benchmark DAPM, we set out to measure the wealth-consumption ratio without committing to a fully-specified equilibrium model. We use a flexible factor model for the stochastic discount factor (henceforth SDF), familiar from the no-arbitrage term structure literature (Duffie and Kan (1996), Dai and Singleton (2000), and Ang and Piazzesi (2003)), and combine it with a vector auto-regression (VAR) for the dynamics of stock returns, bond yields, and consumption and labor income growth, familiar from the methodology of Campbell (1991, 1993, 1996). Like Ang and Piazzesi (2003), we assume that the log SDF is affine in innovations to the state vector, with market prices of risk that are also affine in the same state vector. In a first step we estimate the VAR dynamics of the state. In a second step, we estimate the market prices of aggregate risk. The latter are pinned down by three sets of moments. The first set matches the time-series of nominal bond yields as well as the Cochrane and Piazzesi (2005) bond risk premium. Yields are affine functions of the state, as shown in Duffie and Kan (1996). The second set matches the time series of the price-dividend ratio on the aggregate stock market as well as the equity risk premium. We also impose the present value model: the stock price is the expected present-discounted value of future dividends. The third set uses a cross-section of equity returns to form factor-mimicking portfolios for consumption growth and for labor income growth; these are the linear combinations of assets that have the highest correlations with consumption and labor income growth, respectively. We match the time-series of expected excess returns on these two factor-mimicking portfolios. Our SDF model is flexible enough to provide a close fit for the risk premia on bonds and stocks. With the prices of aggregate risk inferred from traded assets, we price the claims to aggregate consumption and aggregate labor income. This approach has two advantages. First, it avoids making somewhat arbitrary assumptions about the expected rate of return (discount rate) on human wealth, which is unobserved. Second, it avoids using data on housing, durable, and private business wealth, which are often measured at book values and with substantial error. Instead, we only use frequently-traded, precisely-measured stock and bond price data and infer the conditional market prices of aggregate risk from them.

In the benchmark model, we assume that stock and bond prices capture all sources of aggregate risk. This spanning condition is naturally satisfied in standard dynamic asset pricing models, where all aggregate shocks affect the stochastic discount factor and hence asset prices. To guard against the possibility that this condition is not satisfied in the data, we compute an upper bound on the non-traded consumption risk premium. We do so by ruling out good deals, following Cochrane and Saa-Requejo (2000), and we show that there is not enough non-traded consumption risk to alter our conclusion that the consumption risk premium is substantially below the equity risk premium. The validity of our measurement does not rely on market completeness or on the tradeability of human wealth. The approach remains valid in a world with uninsurable labor income risk, in the presence of generic borrowing or wealth constraints, and even if most households only trade in a risk-free asset. If a subset of households has access to the stock and bond markets, the SDF that prices stocks and bonds also prices the consumption and labor income stream.

The low consumption risk premium and the associated high wealth-consumption ratio imply that US households have more wealth than one might think. Our estimates imply that the average household had $3 million of total wealth in 2006. The dynamics of the wealth-consumption ratio are largely driven by the dynamics of real bond yields. As a result, we find that between 1979 and 1981 when real interest rates rose, $533,000 of per capita wealth (in 2006 dollars) was destroyed. Afterwards, as real yields fell, real per capita wealth increased without interruption from $790,000 in 1981 to $3 million in 2006. Greenspan recently argued that the run-up in housing markets in more than two dozen countries between 2001 and 2006 was most likely caused by the decline in real long-term interest rates (Financial Times, April 6, 2008). Our evidence supports his hypothesis for all of US household wealth. Moreover, the timing of the 1979-81 wealth destruction did not coincide with the stock market crash of 1973-74. Likewise, total wealth was hardly affected by the spectacular decline in the stock market that started in 2000. The correlation between realized total wealth returns and stock returns is only 0.12, while the correlation with realized 5-year government bond returns is .50. This suggests that most of the variation in the investment opportunity set of the average US household comes from changes in interest rates not in risk premia on stocks.

On average, the risk-return properties of human wealth closely resemble those of total wealth. We estimate human wealth to be 90 percent of total wealth. This estimate is consistent with Jorgenson and Fraumeni (1989), whose calculations also suggest a 90 percent human wealth share. We estimate that the average household had about $2.6 million in human wealth in 2006. While this number may seem large at first, it pertains to an infinitely-lived household. The value of the first 35 years of aggregate labor income is $840,000. The other two-thirds represent the value of the labor income claim of future generations. The $840,000 amount corresponds to an annuity income of $27,800, close to per capita labor income data in 2006. This $840,000 human wealth number is twelve times higher than the per capita value of residential real estate wealth. This multiple is up from a value of ten in 1981, implying that human wealth grew even faster than housing wealth over the last twenty-five years.

Finally, we compare our results to the predictions of either the simplest or the best leading DAPMs. First, the simple Gordon growth model implies an average wealth-consumption ratio very close to the one we estimate. The discount rate on the consumption claim is 3.49% per year (a consumption risk premium of 2.17% plus a risk-free rate of 1.74% minus a Jensen term of 0.42%); its cash-flow growth rate is 2.34%: $87 = 1/(.0349 - .0234)$. Of course, this calculation ignores the interesting dynamics of the wealth-consumption ratio. Second, we show that two of the leading DAPMs, the long-run risk model of Bansal and Yaron (2004) and the external habit model of Campbell and Cochrane (1999), have very different predictions for the properties of the wealth-consumption ratio, even though they match the same moments of stock returns. Our goal is not to statistically test these models, since our estimation procedure does not nest them, but simply to highlight the key role of the wealth-consumption ratio. Interestingly, the external habit model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and Yaron (2004) have quite different predictions for the wealth-consumption ratio and total wealth returns.[1,2] The long-run risk (LRR) model generates the observed difference between the risk-return characteristics of equity and total wealth because equity (dividends) is more exposed to long-run cash flow risk than total wealth (consumption). It generates a much lower and less volatile wealth-consumption ratio than the price-dividend ratio on equity. The average wealth-consumption ratio in the benchmark LRR model is 87, the same value we estimate in the data, which shows that our numbers are consistent with a standard equilibrium asset pricing model. The external habit (EH) model has an average wealth-consumption ratio of only 12. The low wealth-consumption ratio and associated high consumption risk premium arise because the consumption claim and equity have very similar risk characteristics in this model. In addition, there are some interesting differences between the two models' predictability properties for total wealth returns.

Our approach is closely related to earlier work by Bekaert, Engstrom, and Grenadier (2005), who combine features of the LRR and EH model into an affine pricing model that is calibrated to match moments of stock and bond returns. In contemporaneous work, Lettau and Wachter (2007) also match moments in stock and bond markets with an affine model, while Campbell, Sunderam, and Viceira (2007) study time-varying correlations between bond and stock returns in a quadratic

[1] Early contributions in the habit literature include Abel (1990), Constantinides (1990), Ferson and Constantinides (1991), Abel (1999). See Menzly, Santos, and Veronesi (2004) and Wachter (2006) for more recent contributions. Verdelhan (2007) explores the international finance implications and Chen, Collin-Dufresne, and Goldstein (2008) the implications for credit spreads. Chen and Ludvigson (2007) estimate the habit process for a class of EH models.

[2] Hansen, Heaton, and Li (2008), Parker and Julliard (2005) and Malloy, Moskowitz, and Vissing-Jorgensen (2009) measure long-run risk based on leads and long-run impulse responses of consumption growth. Bansal, Kiku, and Yaron (2006) estimate the long-run risk model. Piazzesi and Schneider (2006) study its implications for the yield curve, Bansal, Dittmar, and Lundblad (2005) study the implications for the cross-section of equity portfolios, Benzoni, Collin-Dufresne, and Goldstein (2008) for options, and Colacito and Croce (2005) and Bansal and Shaliastovich (2007) for international finance. Chen, Favilukis, and Ludvigson (2008) estimate a model with recursive preferences, while Bansal, Gallant, and Tauchen (2007) estimate both long-run risk and external habit models, and Yu (2007) compares correlations between consumption growth and stock returns across the two models.

framework. The focus of our work is on measuring the wealth-consumption (wc) ratio. Lettau and Ludvigson (2001a, 2001b) measure the cointegration residual between log consumption, broadly-defined financial wealth, and labor income, *cay*. The construction of *cay* assumes a constant price-dividend ratio on human wealth. Therefore, human wealth does not contribute to the volatility of the wealth-consumption ratio. Also, *cay* uses as an input the aggregate household wealth data that we try to avoid because of the measurement issues mentioned above. Shiller (1995), Campbell (1996), and Jagannathan and Wang (1996) make assumptions about the properties of expected human wealth returns which are not born out by our estimation exercise. Lustig and Van Nieuwerburgh (2007b) back out the properties of human wealth returns that are consistent with observed consumption growth in the context of the LRR model. Finally, Alvarez and Jermann (2004) estimate the consumption risk premium in order to back out the cost of consumption fluctuations from asset prices. Their log SDF is linear in aggregate consumption growth and the market return. Their model is calibrated to match the unconditional equity premium; it does not allow for time-varying risk premia. They estimate a smaller consumption risk premium of 0.2 percent, and hence a much higher average wealth-consumption ratio. We show that allowing for time-variation in risk premia and matching conditional moments of bond and stock returns raises the estimated consumption risk premium by 2 percent and lowers the wealth-consumption ratio substantially.

We start by measuring the wealth-consumption ratio in the data. Section 1 describes the state variables and their law of motion, while Section 2 shows how we pin down the risk price parameters. Section 3 then describes the estimation results. Section 4 shows that the wealth-consumption ratio estimates are robust to different specifications of the state variables. Section 5 studies the properties of the wealth-consumption ratio in the LRR and EH models.

1 Measuring the Wealth-Consumption Ratio in the Data

Our objective is to estimate the wealth-consumption ratio and the return on total wealth, defined in Section 1.1. Section 1.2 argues that this can be done with a minimal set of assumptions. Section 1.3 describes the state variables and their VAR dynamics.

1.1 Definitions

We start from the aggregate budget constraint:

$$W_{t+1} = R^c_{t+1}(W_t - C_t). \tag{1}$$

The beginning-of-period (or cum-dividend) total wealth W_t that is not spent on aggregate consumption C_t earns a gross return R^c_{t+1} and leads to beginning-of-next-period total wealth W_{t+1}. The return on a claim to aggregate consumption, the *total wealth return*, can be written as

$$R^c_{t+1} = \frac{W_{t+1}}{W_t - C_t} = \frac{C_{t+1}}{C_t}\frac{WC_{t+1}}{WC_t - 1}.$$

Aggregate consumption is the sum of non-durable and services consumption, which includes housing services consumption, and durable consumption. In what follows, we use lower-case letters to denote natural logarithms. We start by using the Campbell (1991) approximation of the log total wealth return $r^c_t = \log(R^c_t)$ around the long-run average log wealth-consumption ratio $A^c_0 \equiv E[w_t - c_t]$.[3]

$$r^c_{t+1} = \kappa^c_0 + \Delta c_{t+1} + wc_{t+1} - \kappa^c_1 wc_t, \tag{2}$$

where we define the log *wealth-consumption* ratio wc as

$$wc_t \equiv \log\left(\frac{W_t}{C_t}\right) = w_t - c_t.$$

The linearization constants κ^c_0 and κ^c_1 are non-linear functions of the unconditional mean wealth-consumption ratio A^c_0:

$$\kappa^c_1 = \frac{e^{A^c_0}}{e^{A^c_0} - 1} > 1 \quad \text{and} \quad \kappa^c_0 = -\log\left(e^{A^c_0} - 1\right) + \frac{e^{A^c_0}}{e^{A^c_0} - 1} A^c_0. \tag{3}$$

1.2 Valuing Human Wealth

The total wealth portfolio includes human wealth. An important question is under what assumptions one can measure the returns on human wealth, and by extension on total wealth, from the returns on traded assets like bonds and stocks. The most direct way to derive the aggregate budget constraint in (1) is by assuming that the representative agent can trade all wealth, including her human wealth. Starting with Campbell (1993), the literature has made this assumption explicitly. In reality, households cannot directly trade claims on their labor income, and the securities they do trade do not fully hedge idiosyncratic labor income risk. They also bear idiosyncratic risk in the form of housing wealth or private business wealth. Finally, a substantial fraction of households do not participate in the stock market but only own a bank account. We argue that the tradeability assumption on human wealth is not necessary. Our measurement of total wealth is valid in a setting with heterogeneous agents who face non-tradeable, non-insurable labor income risk, as well as potentially binding borrowing constraints. Appendix A contains the heterogeneous agent model while we report here the general argument.

[3]Throughout, variables with a subscript zero denote unconditional averages.

Let aggregate consumption $C_t(z^t)$ and aggregate labor income $L_t(z^t)$ depend on the history of the aggregate shocks $z \in Z$, $z^t = \{z_0, z_1, \cdots, z_t\}$. Households not only face aggregate shocks z, but also idiosyncratic shocks which affect their labor income share of the aggregate endowment. For most of the paper, we make two key assumptions. First, we assume that the traded asset payoffs span the aggregate shocks, but not the idiosyncratic shocks. Second, we assume that free portfolio formation and the law of one price hold. This second assumption implies that there exists a unique SDF in the payoff space (Cochrane 2001). We relax our first assumption in Section 4.1 to study the robustness of our results. The traded asset space is:

$$\underline{X}_t = \mathcal{R}^{Z \times t}$$

We take as our stochastic discount factor (SDF) the projection of any candidate SDF M_t on the space of traded payoffs:

$$\frac{\Pi_t^*}{\Pi_{t-1}^*} = proj\left(M_t | \underline{X}_t\right).$$

This SDF is unique in the payoff space. We let P_t^i be the arbitrage-free price of an asset i with non-negative stochastic payoffs $\{D_t^i\}$ that are measurable with respect to z^t:

$$P_t^i = E_t \sum_{\tau=t}^{\infty} \frac{\Pi_\tau^*}{\Pi_t^*} D_\tau^i. \tag{4}$$

Proposition 1. *In a generic incomplete markets economy with market segmentation and spanning of aggregate shocks, the projection of the SDF on the space of traded payoffs can be used to value a claim to aggregate labor income and a claim to aggregate consumption:*

$$P_t^L = E_0\left[\sum_{t=0}^{\infty} \frac{\Pi_t^*}{\Pi_0^*} L_t(z^t)\right],\ W_t = E_0\left[\sum_{t=0}^{\infty} \frac{\Pi_t^*}{\Pi_0^*} C_t(z^t)\right].$$

The resulting prices are human wealth and total wealth, respectively.

This result follows because aggregate consumption and aggregate labor income only depend on aggregate shocks and hence belong to the space of traded payoffs, given the spanning assumption. We note that this pricing result does not apply to household consumption and household labor income, which contain idiosyncratic shocks. The spanning assumption implies that the part of measured aggregate consumption and labor income that is orthogonal to the traded payoffs is measurement error and is not priced:

$$E_t\left[\left(C_{t+1} - proj\left(C_{t+1}|\underline{X}_{t+1}\right)\right)\Pi_{t+1}^*\right] = 0,\ E_t\left[\left(L_{t+1} - proj\left(L_{t+1}|\underline{X}_{t+1}\right)\right)\Pi_{t+1}^*\right] = 0,$$

where $\underline{X}_{t+1}$ includes the risk-free asset and hence the measurement error is mean zero. A key implication of Proposition 1 is that there cannot be a missing risk factor that only appears in the valuation of non-traded assets, and not in the value of traded assets.

The appendix proves this proposition for an environment where heterogeneous agents face labor income risk, which they cannot trade away because of market incompleteness. We can allow some of these households to trade only a limited menu of assets. For example, they could just have access to a one-period bond. As long as there exists a non-zero set of households who trade in securities that are contingent on the aggregate state of the economy (stocks and long-term bonds) and in the one-period bond, we can (i) recover the aggregate budget constraint in equation (1) from the household budget constraints, and (ii) the claim to aggregate labor income and consumption is priced off the same SDF that prices traded assets such as stocks and bonds. In other words, if there exists a SDF that prices stocks and bonds, it also prices aggregate labor income and consumption.

The aggregate risk spanning assumption seems reasonable, especially because it is satisfied in a large class of general equilibrium asset pricing models, including the ones in Section 5. We find it hard to conceive of shocks to aggregate consumption that do not affect prices of any traded assets. For example, recessions or financial crises certainly affect asset prices. Nevertheless, if we relax the spanning-of-aggregate-uncertainty assumption, the part of aggregate consumption and labor income that is orthogonal to traded payoffs, may have a non-zero price. Section 4.1 studies *good-deal bounds* on the consumption risk premium. We find that there is not enough non-traded consumption risk to alter our conclusions without violating reasonable good-deal bounds.

1.3 Model

State Vector We assume that the following state vector describes the aggregate dynamics of the economy:

$$z_t = [CP_t, y_t^{\$}(1), \pi_t, y_t^{\$}(20) - y_t^{\$}(1), pd_t^m, r_t^m, r_t^{fmpc}, r_t^{fmpy}, \Delta c_t, \Delta l_t]'.$$

The first four elements represent the bond market variables in the state, the next four represent the stock market variables, the last two variables represent the cash flows. The state contains in order of appearance: the Cochrane and Piazzesi (2005) factor, the nominal short rate (yield on a 3-month Treasury bill), realized inflation, the spread between the yield on a 5-year Treasury note and a 3-month Treasury bill, the log price-dividend ratio on the CRSP stock market, the real return on the CRSP stock market, the real return on a factor mimicking portfolio for consumption growth, the real return on a factor mimicking portfolio for labor income growth, real per capita consumption growth, and real per capita labor income growth. This state variable is observed at quarterly frequency from 1952.I until 2006.IV (220 observations).[4] Appendix B describes data sources and definitions in detail. All of the variables represent asset prices we want to match or cash flows we need to price (consumption and labor income growth).

The bond risk factor and the factor mimicking portfolios deserve further explanation. Cochrane and Piazzesi (2005) show that a linear combination of forward rates is a powerful predictor of one-year excess bond returns. Following their procedure, we construct 1- through 5-year forward rates from our quarterly nominal yield data, as well as one-year excess returns on 2- through 5-year nominal bonds. We regress the average of the 2- through 5-year excess returns on a constant, the one-year yield, and the 2- through 5-year forward rates. The regression coefficients display a tent-shaped function, very similar to the one reported in Cochrane and Piazzesi (2005). The state variable CP_t is the fitted value of this regression.

Since the aggregate stock market portfolio has a modest 26% correlation with consumption growth, we use additional information from the cross-section of stocks to learn about the consumption and labor income claims. After all, our goal is to price a claim to aggregate consumption and labor income using as much information as possible from traded assets. We use the 25 size- and value-portfolio returns to form a consumption growth factor mimicking portfolio (fmp) and a labor income growth fmp.[5] The consumption (labor income) growth fmp has a 43% (50%) correlation with consumption (labor income) growth. These two fmp returns have a mutual correlation of 70%. The fmp returns are lower on average than the stock return (2.32% and 4.70% versus 7.35% per annum) and are less volatile (6.66% and 13.55% versus 16.68% volatility per annum).[6]

State Evolution Equation We assume that this $N \times 1$ vector of state variables follows a Gaussian VAR with one lag:

$$z_t = \Psi z_{t-1} + \Sigma^{\frac{1}{2}} \varepsilon_t,$$

with $\varepsilon_t \sim i.i.d.\mathcal{N}(0, I)$ and Ψ is a $N \times N$ matrix. The vector z is demeaned. The covariance matrix of the innovations is Σ. We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}\prime}$. $\Sigma^{\frac{1}{2}}$ has non-zero elements only on and below the diagonal. The Cholesky decomposition makes the order of the variables in z important. For example, the innovation to consumption growth is a linear combination of its own (orthogonal) innovation and the innovations to all state variables that precede it. Consumption and labor income growth are placed after the bond and stock variables

[4] Many of these state variables have a long tradition in finance as predictors of stock and bond returns. For example, Ferson and Harvey (1991) study the yield spread, the short rate and consumption growth.

[5] We regress real per capita consumption growth on a constant and the returns on the 25 size and value portfolios (Fama and French 1992). We then form the fmp return series as the product of the 25 estimated loadings and the 25 portfolio return time series. In the estimation, we impose that the fmp weights sum to one and that none of the weights are greater than one in absolute value. We follow the same procedure for the labor income growth fmp.

[6] Interestingly, the same correlation for dividend growth is only 38%. In the estimation, we ensure that our model matches the equity premium. Hence, there is no sense in which the low correlation of consumption growth with returns precludes a high consumption risk premium.

because we use the prices of risk associated with the first eight innovations to value the consumption and labor income claims.

To keep the model parsimonious, we impose additional structure on the companion matrix Ψ. Only the bond market variables -first four- govern the dynamics of the nominal term structure. For example, this structure allows for the CP factor to predict future bond yields, or for the short-term yield and inflation to move together; Ψ_{11} below is a 4×4 matrix of non-zero elements. It also captures that stock returns, the price-dividend ratio on stocks, or the factor-mimicking portfolio returns do not predict future yields; Ψ_{12} is a 4×4 matrix of zeroes. The second block describes the dynamics of the aggregate stock market price-dividend ratio and return, which we assume depends not only on their own lags but also on the the bond market variables. This allows for aggregate stock return predictability by the short rate, the yield spread, inflation, the CP factor, the price dividend-ratio, and lagged aggregate returns, all of which have been shown in the empirical asset pricing literature. The elements Ψ_{21} and Ψ_{22} are 2×4 and 2×2 matrices of non-zero elements. The fmp returns have the same predictability structure as the aggregate stock return, so that Ψ_{31} and Ψ_{32} are 2×4 and 2×2 matrices of non-zero elements. In our benchmark model, consumption and labor income growth do not predict future bond and stock market variables; Ψ_{14}, Ψ_{24}, and Ψ_{34} are all matrices of zeroes. Finally, the VAR structure allows for rich cash flow dynamics: expected consumption growth depends on the first nine state variables and expected labor income growth depends on all lagged state variables; Ψ_{41}, Ψ_{42}, and Ψ_{43} are 2×4, 2×2, and 2×2 matrices of non-zero elements, and Ψ_{44} is a 2×2 matrix with one zero in the upper-right corner.[7] In sum, our benchmark Ψ matrix has the following block-diagonal structure:

$$\Psi = \begin{pmatrix} \Psi_{11} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \Psi_{21} & \Psi_{22} & \mathbf{0} & \mathbf{0} \\ \Psi_{31} & \Psi_{32} & \mathbf{0} & \mathbf{0} \\ \Psi_{41} & \Psi_{42} & \Psi_{43} & \Psi_{44} \end{pmatrix}.$$

In section 4, we explore various alternative restrictions on Ψ. These do not materially alter the dynamics of the estimated wealth-consumption ratio. We estimate Ψ by OLS, equation-by-equation, and we form each innovation as follows $z_{t+1}(\cdot) - \Psi(\cdot, :)z_t$. We compute their (full rank) covariance matrix Σ.

To fix notation, we denote aggregate consumption growth by $\Delta c_t = \mu_c + e_c' z_t$, where μ_c denotes the unconditional mean consumption growth rate and e_c is $N \times 1$ and denotes the column of an $N \times N$ identity matrix that corresponds to the position of Δc in the state vector. The nominal 1-quarter rate is $y_t^{\$}(1) = y_0^{\$}(1) + e_{yn}' z_t$, where $y_0^{\$}(1)$ is the unconditional average nominal short rate

[7]Several of the state variables have been shown to predict consumption growth before. For example, Harvey (1988) finds that expected real interest rates forecast future consumption growth.

and e_{yn} selects the second column of the identity matrix. Likewise, $\pi_t = \pi_0 + e_\pi' z_t$ is the (log) inflation rate between $t-1$ and t with unconditional mean π_0, etc.

Stochastic Discount Factor We adopt a specification of the SDF that is common in the no-arbitrage term structure literature, following Ang and Piazzesi (2003). The nominal pricing kernel $M_{t+1}^{\$} = \exp(m_{t+1}^{\$})$ is conditionally log-normal, where lower case letters continue to denote logs:

$$m_{t+1}^{\$} \quad = \quad -y_t^{\$}(1) - \frac{1}{2}\Lambda_t' \Lambda_t - \Lambda_t' \varepsilon_{t+1}. \tag{5}$$

The real pricing kernel is $M_{t+1} = \exp(m_{t+1}) = \exp(m_{t+1}^{\$} + \pi_{t+1})$.[8] Each of the innovations in the vector ε_{t+1} has its own market price of risk. The $N \times 1$ market price of risk vector Λ_t is assumed to be an affine function of the state:

$$\Lambda_t = \Lambda_0 + \Lambda_1 z_t,$$

for an $N \times 1$ vector Λ_0 and a $N \times N$ matrix Λ_1. The matrix $\Lambda_{1,11}$ contains the bond risk prices, $\Lambda_{1,21}$ and $\Lambda_{1,22}$ contain the aggregate stock risk prices, and $\Lambda_{1,31}$ and $\Lambda_{1,32}$ the fmp risk prices. Importantly, every restriction on Ψ implies a restriction on the elements of the market price of risk we estimate below. Because only bond variables drive the expected returns on bonds, only shocks to the bond variables can affect bond risk premia. For example, the assumption that short term interest rate dynamics do not depend on the price-dividend ratio in the stock market enables us to set the element on the second row and fifth column of Λ_1 equal to zero. Likewise, because the last four variables in the VAR cannot affect expected stock and fmp returns, their (orthogonalized) shocks do not affect risk premia on stocks. Finally, under our assumption of spanning of aggregate uncertainty, the part of the shocks to consumption growth and labor income growth that is orthogonal to the bond and stock innovations is not priced. Thus, $\Lambda_{1,41}$, $\Lambda_{1,42}$, $\Lambda_{1,43}$, and $\Lambda_{1,44}$ are zero matrices. This leads to the following structure for Λ_1:

$$\Lambda_1 = \begin{pmatrix} \Lambda_{1,11} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \Lambda_{1,21} & \Lambda_{1,22} & \mathbf{0} & \mathbf{0} \\ \Lambda_{1,31} & \Lambda_{1,32} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix},$$

We impose corresponding zero restrictions on the mean risk premia in the vector Λ_0: $\Lambda_0 = [\Lambda_{0,1},\ \Lambda_{0,2},\ \Lambda_{0,3}\ \mathbf{0}]'$, where $\Lambda_{0,1}$ is 4×1, and $\Lambda_{0,2}$ and $\Lambda_{0,3}$ are 2×1 vectors. We provide further details on Λ_0 and Λ_1 below. In Section 4.1, we relax the spanning assumption. We derive an

[8]It is also conditionally Gaussian. Note that the consumption-CAPM is a special case of this where $m_{t+1} = \log\beta - \alpha\mu_c - \alpha\eta_{t+1}$ and η_{t+1} denotes the innovation to real consumption growth and α the coefficient of relative risk aversion.

upper bound on the consumption risk premium by increasing the risk price for the consumption growth innovation in $\Lambda_{0,4} > 0$.

The Wealth-Consumption Ratio and Total Wealth Returns In this exponential-Gaussian setting, the log wealth-consumption ratio is an affine function of the state variables:

Proposition 2. *The log wealth-consumption ratio is a linear function of the (demeaned) state vector* z_t

$$wc_t = A_0^c + A_1^{c\prime} z_t,$$

where the mean log wealth-consumption ratio A_0^c *is a scalar and* A_1^c *is the* $N \times 1$ *vector which jointly solve:*

$$0 = \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c - y_0(1) + \frac{1}{2}(e_c + A_1^c)'\Sigma(e_c + A_1^c) - (e_c + A_1^c)'\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right) \quad (6)$$

$$0 = (e_c + e_\pi + A_1^c)'\Psi - \kappa_1^c A_1^{c\prime} - e_{yn}' - (e_c + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_1. \quad (7)$$

In equation (6), $y_0(1)$ denotes the average real one-period bond yield. The proof uses the Euler equation for the (linear approximation of the) total wealth return in equation (2) and is detailed in Appendix C.1. Once we have estimated the market prices of risk Λ_0 and Λ_1 (Section 2), equations (6) and (7) allow us to solve for the mean log wealth-consumption ratio (A_0^c) and its dependence on the state (A_1^c). This is a system of $N+1$ non-linear equations in $N+1$ unknowns; it is non-linear because of equation (3) and can easily be solved numerically.

This solution and the total wealth return definition in (2) imply that the log real total wealth return equals:

$$r_{t+1}^c = r_0^c + \left[(e_c + A_1^c)'\Psi - \kappa_1^c A_1^{c\prime}\right] z_t + (e_c' + A_1^{c\prime})\Sigma^{\frac{1}{2}}\varepsilon_{t+1}, \quad (8)$$

$$r_0^c = \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c. \quad (9)$$

Equation (9) defines the average total wealth return r_0^c. The conditional Euler equation for the total wealth return, $E_t[M_{t+1}R_{t+1}^c] = 1$, implies that the conditional consumption risk premium satisfies:

$$\begin{aligned} E_t\left[r_{t+1}^{c,e}\right] \equiv E_t\left[r_{t+1}^c - y_t(1)\right] + \frac{1}{2}V_t[r_{t+1}^c] &= -Cov_t\left[r_{t+1}^c, m_{t+1}\right] \\ &= (e_c + A_1^c)'\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right) + (e_c + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_1 z_t, \end{aligned} \quad (10)$$

where $E_t\left[r_{t+1}^{c,e}\right]$ denotes the expected log return on total wealth in excess of the real risk-free rate $y_t(1)$, and corrected for a Jensen term. The first term on the last line is the average consumption risk premium (see equation 6). This is a key object of interest which measures how risky total wealth is. The second mean-zero term governs the time variation in the consumption risk premium (see equation 7).

The structure we impose on Ψ and on the market prices of risk is not overly restrictive. A Campbell-Shiller decomposition of the wealth-consumption ratio into an expected future consumption growth component (Δc_t^H) and an expected future total wealth returns component (r_t^H), detailed in Appendix C, delivers the following expressions:

$$\Delta c_t^H = e_c' \Psi (\kappa_1^c I - \Psi)^{-1} z_t \quad \text{and} \quad r_t^H = \left[(e_c + A_1^c)'\Psi - \kappa_1^c A_1^{c\prime}\right] (\kappa_1^c I - \Psi)^{-1} z_t.$$

Despite the restrictions we impose on Ψ and Λ_t, both the cash flow component and the discount rate component depend on all the stock and the bond components of the state. In the case of cash flows, this follows from the fact that expected consumption growth depends on all lagged stock and bond variables in the state. In the case of discount rates, there is additional dependence through A_1^c, which itself is a function of the first nine state variables. The cash flow component does not directly depend on the risk prices (other than through κ_1^c) while the discount rate component depends on all risk prices of stocks and bonds through A_1^c. This flexibility implies that our model can accommodate a large consumption risk premium; when the covariances between consumption growth and the other aggregate shocks are large and/or when the unconditional risk prices in Λ_0 are sufficiently large. In fact, in our estimation, we choose Λ_0 large enough to match the equity premium. A low estimate of the consumption risk premium and hence a high wealth-consumption ratio are not a foregone conclusion.

2 Estimating the Market Prices of Risk

To compute the wealth-consumption ratio we need estimates of the market price of risk parameters. We identify Λ_0 and Λ_1 from the moments of bond yields and stock returns. The estimation proceeds in four stages.

1. In a first step, we estimate the risk prices in the *bond market block* $\Lambda_{0,1}$ and $\Lambda_{1,11}$ by matching the two yields in the state vector. Because of the block diagonal structure, we can estimate these separately.

2. In a second step, we estimate the risk prices in the *stock market block* $\Lambda_{0,2}$, $\Lambda_{1,21}$, and $\Lambda_{1,22}$ jointly with the bond risk prices, taking the estimates from step 1 as starting values.

3. In a third step, we estimate the fmp risk prices in the *factor mimicking portfolio block* $\Lambda_{0,3}$, $\Lambda_{1,31}$, and $\Lambda_{1,32}$ taking as given the bond and stock risk prices.

4. Finally, we impose over-identifying restrictions on the estimation, such as matching additional nominal yields, imposing the present-value relationship for stocks, and imposing a human wealth share between zero and one. We re-estimate all 5 parameters in Λ_0 and all 26 parameters in Λ_1, starting with the estimates from the third step.

In estimating the market price of risk parameters, we impose that the model replicates the time-series of these traded asset prices that are part of the state vector. Such constraints pin down many parameters of the model, but we often need to choose which ones. In selecting the identified parameters, we follow a simple rule: we associate prices of risk with traded assets instead of non-traded variables. The VAR parameter estimates as well as the estimates for the market prices of risk from the last-stage estimation are listed at the end of Appendix C. We now provide more detail on each of these steps.

2.1 Block 1: Bonds

The first four elements in the state, the Cochrane-Piazzesi factor, the nominal 3-month T-bill yield, the inflation rate, and the yield spread (5-year T-bond minus the 3-month T-bill yield), govern the term structure of interest rates. Together they deliver a four-factor term structure model. In contrast to most of the term structure literature, all factors are observable. The price of a τ-period nominal zero-coupon bond satisfies:

$$P_t^{\$}(\tau) = E_t\left[e^{m_{t+1}^{\$}+\log P_{t+1}^{\$}(\tau-1)}\right].$$

This defines a recursion with $P_t^{\$}(0) = 1$. The corresponding bond yield is $y_t^{\$}(\tau) = -\log(P_t^{\$}(\tau))/\tau$. From Ang and Piazzesi (2003), we know that bond yields in this class of models are an affine function of the state: $y_t^{\$}(\tau) = -\frac{A^{\$}(\tau)}{\tau} - \frac{B^{\$}(\tau)'}{\tau}z_t$. Appendix C.3 formally states and proves this result and provides the recursions for $A^{\$}(\tau)$ and $B^{\$}(\tau)$. Given the block-diagonal structure of Λ_1 and Ψ, only the risk prices in $\Lambda_{0,1}$ and $\Lambda_{1,11}$ affect the yield loadings. That is why, in a first step, we can estimate the bond block separately from the stock block. We do so by matching the time series for the slope of the yield curve and the CP risk factor.

First, we impose that the model prices the 1-quarter and the 20-quarter nominal bond correctly. The condition $A^{\$}(1) = -y_0^{\$}(1)$ guarantees that the one-quarter nominal yield is priced correctly on average, and the condition $B^{\$}(1) = -e_{yn}$ guarantees that the nominal short rate dynamics are identical to those in the data. The short rate and the yield spread are in the state, which implies the following expression for the 20-quarter bond yields:

$$y_t^{\$}(20) = y_0^{\$}(20) + (e'_{yn} + e'_{spr})z_t.$$

Matching the 20-quarter yield implies two sets of parameter restrictions:

$$\begin{aligned} \frac{-1}{20} A^{\$}(20) &= y_0^{\$}(20), & (11) \\ \frac{-1}{20} \left(B^{\$}(20)\right)' &= (e_{yn} + e_{spr})'. & (12) \end{aligned}$$

Equation (11) imposes that the model matches the unconditional expectation of the 5-year nominal yield $y_0^{\$}(20)$. This provides one restriction on Λ_0; it identifies its second element. To match the dynamics of the 5-year yield, we need to free up one row in the bond block of the risk price matrix $\Lambda_{1,11}$; we choose to identify the second row in $\Lambda_{1,11}$. We impose the restrictions (11) and (12) by minimizing the summed square distance between model-implied and actual yields.

Second, we match the time-series of the CP risk factor $(CP_0 + e_{cp}' z_t)$ in order to replicate the dynamics of bond risk premia in the data. We follow the exact same procedure to construct the CP factor in the model as in the data, using the model-implied yields to construct forward rates. By matching the mean of the factor in model and data, we can identify one additional element of Λ_0; we choose the fourth element. By matching the dynamics, we can identify four more elements in $\Lambda_{1,11}$, one in each of the first four columns; we identify the fourth row in $\Lambda_{1,11}$. We impose the restriction that the CP factor is equal in model and data by minimizing their summed squared distance. We now have identified two elements (rows) in $\Lambda_{0,1}$ (in $\Lambda_{1,11}$). The first and third elements (rows) in $\Lambda_{0,1}$ (in $\Lambda_{1,11}$) are zero.

2.2 Block 2: Stocks

In the second step, we turn to the estimation of the risk price parameters in $\Lambda_{1,21}$ and $\Lambda_{1,22}$. We do so by imposing that the model prices excess stock returns correctly; we minimize the summed squared distance between VAR- and SDF-implied excess returns:

$$\begin{aligned} E_t^{VAR}[r_{t+1}^{m,e}] &= r_0^m - y_0(1) + \frac{1}{2} e_{rm}' \Sigma e_{rm} + \left((e_{rm} + e_\pi)' \Psi - e_{yn}' \right) z_t, \\ E_t^{SDF}[r_{t+1}^{m,e}] &= e_{rm}' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime} e_\pi \right) + (e_{rm} + e_\pi)' \Sigma^{\frac{1}{2}} \Lambda_1 z_t, \end{aligned}$$

where r_0^m is the unconditional mean stock return and e_{rm} selects the stock return in the VAR. Matching the unconditional equity risk premium in model and data identifies one additional element in Λ_0; we choose the sixth element (the second element of $\Lambda_{0,2}$). Matching the risk premium dynamics allows us to identify the second row in $\Lambda_{1,21}$ (4 elements) and the second row in $\Lambda_{1,22}$ (2 more elements). Choosing to identify the sixth element (row) of Λ_0 (Λ_1) instead of the fifth row is an innocuous choice. But it is more natural to associate the prices of risk with the traded stock return rather than with the non-traded price-dividend ratio. These six elements in $\Lambda_{1,22}$ must all be non-zero because expected returns in the VAR depend on the first six state variables. The first element of $\Lambda_{0,2}$ and the first rows of $\Lambda_{1,21}$ and $\Lambda_{1,22}$) are zero.

2.3 Block 3: Factor Mimicking Portfolios

In addition, we impose that the risk premia on the fmp coincide between the VAR and the SDF model. As is the case for the aggregate stock return, this implies one additional restriction on Λ_0 and N additional restrictions on Λ_1:

$$\begin{aligned} E_t^{VAR}[r_{t+1}^{fmp,e}] &= r_0^{fmp} - y_0(1) + \frac{1}{2}e'_{fmp}\Sigma e_{fmp} + \left((e_{fmp} + e_\pi)'\Psi - e'_{yn}\right) z_t, \\ E_t^{SDF}[r_{t+1}^{fmp,e}] &= e'_{fmp}\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime} e_\pi\right) + (e_{fmp} + e_\pi)'\Sigma^{\frac{1}{2}}\Lambda_1 z_t, \end{aligned}$$

where r_0^{fmp} is the unconditional average fmp return. There are two sets of such restrictions, one set for the consumption growth and one set for the labor income growth fmp. Matching average expected fmp returns and their dynamics identifies both elements of $\Lambda_{0,3}$. Matching the risk premium dynamics allows us to identify both rows of in $\Lambda_{1,31}$ (4 elements) and $\Lambda_{1,32}$ (4 more elements).

2.4 Over-identifying Restrictions

The stock and bond moments described thus far exactly identify the 5 elements of Λ_0 and the 26 elements of Λ_1. In other words, given the structure of Ψ, they are all strictly necessary to match the levels and dynamics of bond yields and stock returns. For theoretical as well as for reasons of fit, we impose several additional constraints. To avoid over-parametrization, we choose not to let these constraints identify additional market price of risk parameters.

Additional Nominal Yields We minimize the squared distance between the observed and model-implied yields on nominal bonds of maturities 1, 3, 10, and 20 years. These additional yields are useful to match the dynamics of long-term yields. This will be important given that the dynamics of the wealth-consumption ratio turn out to be largely driven by long yields. We impose several other restrictions that force the term structure to be well-behaved at long horizons.[9]

[9] We impose that the average nominal and real yields at maturities 200, 500, 1000, and 2500 quarters are positive, that the average nominal yield is above the average real yield at these same maturities, and that the nominal and real yield curves flatten out. The last constraint is imposed by penalizing the algorithm for choosing a 500-200 quarter yield spread that is above 3% per year and a 2500-500 quarter yield spread that is above 2% per year. Together, they guarantee that the infinite sums we have to compute are well-behaved. None of these additional term structure constraints are binding at the optimum.

Price-Dividend Ratio While we imposed that expected excess equity returns coincide between the VAR and the SDF model, we have not yet imposed that the return on stocks reflects cash flow risk in the equity market. To do so, we require that the price-dividend ratio in the model, which is the expected present discounted value of all future dividends, matches the price-dividend ratio in the data, period by period.[10] To calculate the price-dividend ratio on equity, we use the fact that it must equal the sum of the price-dividend ratios on dividend strips of all horizons (Wachter (2005)):

$$\frac{P_t^m}{D_t^m} = e^{pd_t^m} = \sum_{\tau=0}^{\infty} P_t^d(\tau), \tag{13}$$

where $P_t^d(\tau)$ denotes the price of a τ period dividend strip divided by the current dividend. A dividend strip of maturity τ pays 1 unit of dividend at period τ, and nothing in the other periods. The strip's price-dividend ratio satisfies the following recursion:

$$P_t^d(\tau) = E_t\left[e^{m_{t+1}+\Delta d_{t+1}^m + \log\left(P_{t+1}^d(\tau-1)\right)}\right],$$

with $P_t^d(0) = 1$. Aggregate dividend growth Δd^m is obtained from the dynamics of the pd^m ratio and the stock return r^m through the definition of the stock return. Appendix C.4 formally states and proves that the log price-dividend ratios on dividend strips are affine in the state vector: $\log\left(P_t^d(\tau)\right) = A^m(\tau) + B^{m\prime}(\tau)z_t$. It also provides the recursions for $A^m(\tau)$ and $B^m(\tau)$. See Bekaert and Grenadier (2001) for a similar result. Using (13) and the affine structure, we impose the restriction that the price-dividend ratio in the model equals the one in the data by minimizing their summed squared distance. Imposing this constraint not only affects the price of cash flow risk (the sixth row of Λ_t) but also the real term structure of interest rates (the second and fourth rows of Λ_t). Real yields turn out to play a key role in the valuation of real claims such as the claim to real dividends (equity) or the claim to real consumption (total wealth).[11] As such, the price-dividend ratio restriction turns out to be useful in sorting out the decomposition of the nominal term structure into an inflation component and the real term structure.

[10] This constraint is not automatically satisfied from the definition of the stock return: $r_{t+1}^m = \kappa_0^m + \Delta d_{t+1}^m + \kappa_1^m pd_{t+1}^m - pd_t^m$. The VAR implies a model for expected return and the expected log price-dividend ratio dynamics, which implies expected dividend growth dynamics through the definition of a return. These dynamics are different from the ones that would arise if the VAR contained dividend growth and the price-dividend ratio instead. The reason is that the state vector in the first case contains r_t and pd_t^m, while in the second case it contains Δd_t^m and pd_t^m. For the two models to have identical implications for expected returns and expected dividend growth, one would need to include pd_{t-1}^m as an additional state variable. We choose to include returns instead of dividend growth rates because the resulting properties for expected returns and expected dividend growth rates are more desirable. For example, the two series have a positive correlation of 20%, a number similar to what Lettau and Ludvigson (2005) estimate. See Lettau and Van Nieuwerburgh (2007), Ang and Liu (2007), and Binsbergen and Koijen (2007) for an extensive discussion of the present-value constraint.

[11] Appendix C.3 shows that real bond yields $y_t(\tau)$, denoted without the \$ superscript, are also affine in the state, and provides the recursions for the coefficients.

Human Wealth Share The same way we priced a claim to aggregate consumption, we price a claim to aggregate labor income. We impose that the conditional Euler equation for human wealth returns is satisfied and obtain a log price-dividend ratio which is also affine in the state: $pd_t^l = A_0^l + A_1^l z_t$. (See Corollary 5 in Appendix C.1.) By the same token, the conditional risk premium on the labor income claim is given by:

$$E_t\left[r_{t+1}^{l,e}\right] = (e_{\Delta l} + A_1^l)'\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right) + (e_{\Delta l} + A_1^l)'\Sigma^{\frac{1}{2}}\Lambda_1 z_t.$$

We use μ_l to denote unconditional labor income growth and $e_{\Delta l}$ selects labor income growth in the VAR. We also impose that aggregate labor income grows at the same rate as aggregate consumption ($\mu_l = \mu_c$).[12] We define the labor income share, lis_t, as the ratio of aggregate labor income to aggregate consumption. It is 0.855 on average in our sample. The human wealth share is the ratio of human wealth to total wealth; it is a function of the labor income share and the price-dividend ratios on human and total wealth:

$$hws_t = lis_t \frac{e^{pd_t^l} - 1}{e^{wc_t} - 1}.$$

We impose on the estimation that hws_t lies between 0 and 1 at each time t. At the optimum, this constraint is satisfied.

3 Results

Before studying the estimation results for the wealth-consumption ratio, we check that the model does an adequate job describing the dynamics of the bond yields and of stock returns.

3.1 Model Fit for Bonds and Stocks

The model fits the nominal term structure of interest rates reasonably well. We match the 3-month yield exactly . The first two panels of Figure 1 plot the observed and model-implied *average* nominal yield curve, while Figure 2 plots the entire time-series for the 1-quarter, 1-, 3-, 5-, 10-, and 20-year yields. For the 5-year yield, which is part of the state vector, the average pricing error is -5 basis points (bp) per year. The annualized standard deviation of the pricing error is only 13 bp, and the root mean squared error (RMSE) is 26 bp. For the other four yields, the mean annual

[12] We rescale the level of consumption to end up with the same average labor income share (after imposing $\mu_l = \mu_c$) as in the data (before rescaling). This transformation does not affect growth rates. The assumption is meant to capture that labor income and consumption cannot diverge in the long run. In Section 4, we estimate a model where we impose cointegration between consumption and labor income by including the log consumption-labor income ratio $c - l$ ratio in place of Δl in the state vector. As explained below, we impose that the human wealth share stays between 0 and 1 in all our estimations.

pricing errors range from -18 bp to +61 bp, the volatility of the pricing errors range from 10 to 60 bp, and the RMSE from 24 to 134 bp.[13] While these pricing errors are somewhat higher than the ones produced by term-structure models, our model with only 8 parameters in the term structure block of Λ_1 and no latent variables does a good job capturing the level and dynamics of long yields. Furthermore, most of the term structure literature prices yields of maturities up to 5 years, while we also price the 10-year and 20-year yields, because these matter for pricing long-duration assets. On the dynamics, the annual volatility of the nominal yield on the 5-year bond is 1.36% in the data and 1.29% in the model.

[Figure 1 about here.]

[Figure 2 about here.]

The model also does a good job capturing the bond risk premium dynamics. The right panel of Figure 3 shows a close fit between the Cochrane-Piazzesi factor in model and data. It is a measure of the 1-quarter nominal bond risk premium. The left panel shows the 5-year nominal bond risk premium, defined as the difference between the 5-year yield and the average expected future short term yield averaged over the next 5 years. This long-term measure of the bond risk premium is also matched closely by the model, in large part due to the fact that the long-term and short-term bond risk premia have a correlation of 90%.

[Figure 3 about here.]

The model also manages to capture the dynamics of stock returns quite well. The bottom panel of Figure 4 shows that the model matches the equity risk premium that arises from the VAR structure. The average equity risk premium (including Jensen term) is 6.90% per annum in the data, and 7.06% in the model. Its annual volatility is 9.54% in the data and 9.62% the model. The top panel shows the dynamics of the price-dividend ratio on the stock market. The model, where the price-dividend ratio reflects the present discounted value of future dividends, replicates the price-dividend ratio in the data quarter by quarter.

[Figure 4 about here.]

As in Ang, Bekaert, and Wei (2008), the long-term nominal risk premium on a 5-year bond is the sum of a real rate risk premium (defined the same way for real bonds as for nominal bonds) and the inflation risk premium. The right panel of Figure 5 decomposes this long-term bond risk premium (solid line) into a real rate risk premium (dashed line) and an inflation risk premium

[13]Note that the largest errors occur on the 20-year yield, which is unavailable between 1986.IV and 1993.II. The standard deviation and RMSE on the 10-year yield are only half as big as on the 20-year yield.

(dotted line). The real rate risk premium becomes gradually more important at longer horizons. The left panel of Figure 5 decomposes the 5-year yield into the real 5-year yield (which itself consists of the expected real short rate plus the real rate risk premium), expected inflation over the next 5-years, and the 5-year inflation risk premium. The inflationary period in the late 1970s-early 1980s was accompanied by high inflation expectations and an increase in the inflation risk premium, but also by a substantial increase in the 5-year real yield.[14] Separately identifying real rate risk and inflation risk based on term structure data alone is challenging.[15] We do not have long enough data for real bond yields, but stocks are real assets that contain information about the term structure of real rates. They can help with the identification. For example, higher long real yields in the late 1970s-early 1980s lower the price-dividend ratio on stocks, which indeed was low in the late 1970s-early 1980s (top panel of Figure 4). In terms of average real yields, the third panel of Figure 1 shows yields ranging from 1.74% per year for 1-quarter real bonds to 2.70% per year for 20-year real bonds.

[Figure 5 about here.]

Finally, the model matches the expected returns on the consumption and labor income growth factor mimicking portfolios (fmp) very well. The figure is omitted for brevity. The annual risk premium on the consumption growth fmp is 0.79% with a volatility of 1.67 in data and model. Likewise, the risk premium on the labor income growth fmp is 3.87% in data and model, with volatilities of 1.92 and 1.98%.

3.2 The Wealth-Consumption Ratio

With the estimates for Λ_0 and Λ_1 in hand, it is straightforward to use Proposition 2 and solve for A_0^c and A_1^c from equations (6)-(7). The last column of Table 1 summarizes the key moments of the log wealth-consumption ratio. The numbers in parentheses are small sample bootstrap standard errors, computed using the procedure described in Appendix C.7. We can directly compare the moments of the wealth-consumption ratio with those of the price-dividend ratio on equity. The wc ratio has a volatility of 17% in the data, considerably lower than the 27% volatility of the pd^m ratio. The wc ratio in the data is a persistent process; its 1-quarter (4-quarter) serial correlation is .96 (.85). This is similar to the .95 (.78) serial correlation of pd^m. The volatility of changes in

[14] Inflation expectations in our VAR model have a correlation of 80% with inflation expectations from the Survey of Professional Forecasters (SPF) over the common sample 1981-2006. The 1-quarter ahead inflation forecast error series for the SPF and the VAR have a correlation of 68%. Realized inflation fell sharply in the first quarter of 1981. Neither the professional forecasters nor the VAR anticipated this decline, leading to a high realized real yield. The VAR expectations caught up more quickly than the SPF expectations, but by the end of 1981, both inflation expectations were identical.

[15] Many standard term structure models have a likelihood function with two local maxima with respect to the persistence parameters of expected inflation and the real rate.

the wealth consumption ratio is 4.86%, and because of the low volatility of aggregate consumption growth changes, this translates into a volatility of the total wealth return on the same order of magnitude (4.93%). The corresponding annual volatility of 9.8% is much lower than the 16.7% volatility of stock returns. The change in the wc ratio and the total wealth return have weak autocorrelation (-.11 and -.01 at the 1 and 4 quarter horizons for both), suggesting that total wealth returns are hard to forecast by their own lags. The correlation between the total wealth return and consumption growth is only mildly positive (.19). How risky is total wealth in the data? According to our estimation, the consumption risk premium (calculated from equation 10) is 54 basis points per quarter or 2.17% per year. This results in a mean wealth-consumption ratio (A_0^c) of 5.86 in logs, or 87 in annual levels ($\exp\{A_0^c - log(4)\}$). The consumption risk premium is only one-third as big as the equity risk premium of 6.9%. Correspondingly, the wealth-consumption ratio is much higher than the price-dividend ratio on equity: 87 versus 27. A simple back-of-the-envelope Gordon growth model calculation sheds light on the level of the wealth-consumption ratio. The discount rate on the consumption claim is 3.49% per year (a consumption risk premium of 2.17% plus a risk-free rate of 1.74% minus a Jensen term of 0.42%) and its cash-flow growth rate is 2.34%: $87 = 1/(.0349 - .0234)$. Finally, the volatility of the consumption risk premium is 3.3% per year, one-third of the volatility of the equity risk premium. The standard errors on the moments of the wealth-consumption ratio or total wealth return are sufficiently small so that the corresponding moments of the price-dividend ratio or stock returns are outside the 95% confidence interval of the former. The main conclusion of our measurement exercise is that total wealth is (economically and statistically) significantly less risky than equity.

[Table 1 about here.]

Figure 6 plots the time-series for the annual wealth-consumption ratio, expressed in levels. Its dynamics are to a large extent inversely related to the long real yield dynamics (dashed line in the left panel of Figure 5). For example, the 5-year real yield increases from 2.7% per annum in 1979.I to 7.3% in 1981.III while the wealth-consumption ratio falls from 77 to 46. This corresponds to a loss of $533,000 in real per capita wealth in 2006 dollars.[16] Similarly, the low-frequency decline of the real yield in the twenty-five years after 1981 corresponds to a gradual rise in the wealth-consumption ratio. One striking way to see that total wealth behaves differently from equity is to study it during periods of large stock market declines. During the periods 1973.III-1974.IV and 2000.I-2002.IV, for example, the change in US households' real per capita stock market wealth, including mutual fund holdings, was -46% and -61%, respectively. In contrast, real per capita total wealth changed by -12% and +27%, respectively. Over the full sample, the total wealth return has a correlation of only 12% with the value-weighted real CRSP stock return, while it has a correlation

[16] Real per capita wealth is the product of the wealth-consumption ratio and observed real per capita consumption.

of 46% with realized one-year holding period returns on the 5-year nominal government bond.[17] Likewise, the (1-period) consumption risk premium has a correlation of 22% with the (1-period) equity risk premium, but 66% with the (1-period) nominal bond risk premium.

[Figure 6 about here.]

To show more formally that the consumption claim behaves like a real bond, we compute the discount rate that makes the current wealth-consumption ratio equal to the expected present discounted value of future consumption growth (See Appendix C.2 for details). This is the solid line measured against the left axis of Figure 7. Similarly, we calculate a time series for the discount rate on the dividend claim, the dotted line measured against the right axis. For comparison, we plot the yield on a long-term real bond (50-year) as the dashed line against the right axis. The correlation between the consumption discount rate and the real yield is 99%, whereas the correlation of the dividend discount rate and the real yield is only 44%. In addition, the consumption and dividend discount rates only have a correlation of 47%, reinforcing our conclusion that the data suggest a big divergence between the perceived riskiness of a claim to consumption and a claim to dividends in securities markets.

[Figure 7 about here.]

Consumption Strips A different way of showing that the consumption claim is bond-like is to study yields on consumption strips. Just as the price-dividend ratio on stocks equals the sum of the price-dividend ratios on dividend strips of all maturities, so is the wealth-consumption ratio equal to the price-dividend ratio on all consumption strips. A consumption strip of maturity τ pays 1 unit of consumption at period τ, and nothing in the other periods. It is useful to decompose the yield on the period-τ strip in two pieces. The first component is the yield on a security that pays a certain cash flow $(1+\mu_c)^{\tau}$. The underlying security is a real perpetuity with a certain cash flow which grows at a deterministic consumption growth rate μ_c. The second component is the yield on a security that pays off $C_{\tau}/C_0 - (1+\mu_c)^{\tau}$. It captures pure consumption cash flow risk. Appendix C.5 shows that the log price-dividend ratios on the consumption strips are affine in the state, and details how to compute the yield on its two components. Figure 8 makes clear that the consumption strip yields are mostly comprised of a compensation for time value of money, not consumption cash flow risk.

[Figure 8 about here.]

[17] A similarly low correlation of 12% is found between total wealth returns and the Flow of Fund's measure of the growth rate in real per capita household net worth, a broad measure of financial wealth. The correlation of the total wealth return with the Flow of Fund's growth rate of real per capita housing wealth is 0.11.

The consumption strips with and without cash-flow risk are also useful to compute the welfare cost of consumption fluctuations. Alvarez and Jermann (2004) show that the marginal cost of such fluctuations equals the ratio of the price of a claim to consumption without cash-flow risk (growing at a deterministic rate μ_c) to the price of a claim to consumption with cash-flow risk minus one. Our estimates imply a marginal cost of consumption fluctuations of 37% on average, or about twice as high as the benchmark estimate in Alvarez and Jermann (2004). Because consumption risk premia fluctuate over time in our model, so does the cost of consumption fluctuations: between 28 and 46%. About 40% of that marginal cost of all consumption fluctuations is the cost of business-cycle frequency fluctuations (Alvarez and Jermann 2004).

3.3 Human Wealth Returns

Our estimates indicate that the bulk of total wealth is human wealth. The human wealth share fluctuates between 85 and 96%, with an average of 90%. Interestingly, Jorgenson and Fraumeni (1989) also calculate a 90% human wealth share. The average price-dividend ratios on human wealth is slightly above the one on total wealth (94 versus 87 in annual levels). The risk premium on human wealth is very similar to the one for total wealth (2.19 versus 2.17% per year). The price-dividend ratios and risk premia on human wealth and total wealth have a 99% correlation. In line with the findings of Lustig and Van Nieuwerburgh (2007b), we estimate only a weak contemporaneous correlation between risk premia on human wealth and on equity (0.19).

Existing approaches to measuring total wealth make ad hoc assumptions about expected human wealth returns. The model of Campbell (1996) assumes that expected human wealth returns are equal to expected returns on financial assets. This is a natural benchmark when financial wealth is a claim to a constant fraction of aggregate consumption. Shiller (1995) models a constant discount rate on human wealth. Jagannathan and Wang (1996) assume that expected returns on human wealth equal the expected labor income growth rate; the resulting price-dividend ratio on human wealth is constant. The construction of *cay* in Lettau and Ludvigson (2001a) makes that same assumption. These models can be thought of as special case of ours, imposing additional restrictions on the market prices of risk Λ_0 and Λ_1. Our estimation results indicate that expected excess human wealth returns have an annual volatility of 3.7%. This is substantially higher than the volatility of expected labor income growth (0.7%), but much lower than that of the expected excess returns on equity (9.6%). Lastly, average (real) human wealth returns (3.9%) are much lower than (real) equity returns (7.4%), but higher than (real) labor income growth (2.3%) and the (real) short rate (1.7%). In sum, our approach avoids having to make arbitrary assumptions on unobserved human wealth returns. Our findings do not quite fit any of the assumptions on human wealth returns made in previous work.

How much wealth, and in particular human wealth, do our estimates imply? In real 2006 dollars, total per capita wealth increased from \$1 million to \$3 million between 1952 and 2006. The thick solid line in the left panel of Figure 9 shows the time series. Of this, \$2.6 million was human wealth in 2006 (dashed line), while the remainder is non-human wealth (dotted line, plotted in the right panel). To better judge whether this is a realistic number, we compute what fraction of human wealth accrues in the first 35 years. This fraction is the price of the first 140 quarterly labor income strips divided by the price of all labor income strips. The labor income strip prices are computed just like the consumption strip prices. On average, 33% of human wealth pertains to the first 35 years. In 2006, this implies a human wealth value of \$840,000 per capita (thin solid line in right panel). This amount is the price of a 35-year annuity with a cash flow of \$27,850 which grows at the average labor income growth rate of 2.34% and is discounted at the average real rate of return on human wealth of 3.41%. This model-implied annual income of \$27,850 is close to the \$25,360 US per capita labor income at the end of 2006 (National Income and Products Accounts, Table 2.1). To further put this number in perspective, we compare the "first 35 years" human wealth number to the per capita value of residential real estate wealth from the Flow of Funds. It is 12.3 times higher than real estate wealth in 2006. This multiple is up from a value of 9.7 in 1981.III, so that human wealth grew even faster than housing wealth over the last twenty-five years. In sum, human wealth has been an important driver behind the fast wealth accumulation.

[Figure 9 about here.]

Finally, we compare non-human wealth, the difference between our estimates for total and for human wealth, with the Flow of Funds series for household net worth. The latter is the sum of equity, bonds, housing wealth, durable wealth, private business wealth, and pension and life insurance wealth minus mortgage and credit card debt. Our non-human wealth series is on average 1.7 times the Flow of Funds series. This ratio varies over time: it is 2.2 at the beginning and at the end of the sample, and it reaches a low of 0.7 in 1973. We chose not to use the Flow of Funds net worth data in our estimation because many of the wealth categories are hard to measure accurately or are valued at book value (e.g., private business wealth). Arguably, only the equity component for publicly traded companies is measured precisely, and this may explain why the dynamics of the household net worth series are to a large extent driven by variation in stock prices (Lettau and Ludvigson (2001a)).[18] It is reassuring that our non-human wealth measure exceeds the net worth series. After all, our series measures the present discounted value of all future non-labor income. This includes the value of growth options that will accrue to firms that have not been born yet, the same way human wealth includes labor income from future generations.

[18] Lettau and Ludvigson (2001a)'s measure $-cay$ falls during the stock market crashes of 1974 and 2000-02. It has a correlation of only 0.16 with our wealth-consumption measure while it has a correlation of 0.37 with the price-dividend ratio on stocks.

3.4 Predictability Properties

Our analysis so far has focused on unconditional moments of the total wealth return. The conditional moments of total wealth returns are also very different from those of equity returns. The familiar Campbell and Shiller (1988) decomposition for the wealth-consumption ratio shows that the wealth-consumption ratio fluctuates either because it predicts future consumption growth rates (Δc_t^H) or because it predicts future total wealth returns (r_t^H):

$$V\left[wc_t\right] = Cov\left[wc_t, \Delta c_t^H\right] + Cov\left[wc_t, -r_t^H\right] = V\left[\Delta c_t^H\right] + V\left[r_t^H\right] - 2Cov\left[r_t^H, \Delta c_t^H\right].$$

The second equality suggests an alternative decomposition into the variance of expected future consumption growth, expected future returns, and their covariance. Finally, it is straightforward to break up $Cov\left[wc_t, r_t^H\right]$ into a piece that measures the predictability of future excess returns, and a piece that measures the covariance of wc_t with future risk-free rates. Our no-arbitrage methodology delivers analytical expressions for all variance and covariance terms (See Appendix C.2).

We draw three main empirical conclusions. First, the mild variability of the wc ratio implies only mild (total wealth) return predictability. This is in contrast with the high variability of pd^m. Second, 98.4% of the variability in wc is due to covariation with future total wealth returns while the remaining 1.6% is due to covariation with future consumption growth. Hence, the wealth-consumption ratio predicts future returns (discount rates), not future consumption growth rates (cash flows). Using the second variance decomposition, the variability of future returns is 97%, the variability of future consumption growth is 0.3% and their covariance is 2.7% of the total variance of wc. This variance decomposition is similar to the one for equity. Third, 69.6% of the 98.4% covariance with returns is due to covariance with future *risk-free rates*, and the remaining 28.7% is due to covariance with future *excess* returns. The wealth-consumption ratio therefore mostly predicts future variation in interest rates, not in risk premia. The exact opposite holds for equity: the bulk of the predictability of the pd^m ratio for future stock returns is predictability of excess returns (74.7% out of 97.0%). In sum, the conditional asset pricing moments also reveal interesting differences between equity and total wealth. Again, they point to the link between the consumption claim return and interest rates.

Finally, these results do not change if we allow for leverage following Abel (1999)'s approach. We price a claim to C_t^κ, with a leverage ratio of $\kappa = 3$. The average price-dividend ratio on this claim (59) is still much higher than that on equity (27). In addition, the variance decomposition of the price-dividend ratio on that levered consumption claim is virtually identical to that of the wealth-consumption ratio (not reported in the table). In particular, the covariance with future risk-free rates still accounts for 58% of the variance, while only 36% is accounted for by the covariance with excess returns. In sum, introducing leverage does not close the gap between the consumption claim and the equity claim decomposition.

4 Robustness Analysis

We now consider several robustness tests to our findings. We first relax our spanning assumption. We then entertain several alternative specifications of the dynamics of the state vector. Finally, we consider an estimation at annual frequency.

4.1 Non-Traded Consumption Risk

Sofar we have assumed that all aggregate shocks are spanned by stock and bond prices. This assumption is satisfied in the DAPMs of Section 5. Even in incomplete markets models, asset prices will reflect changes in the income or wealth distribution (e.g., Constantinides and Duffie (1996)). Nevertheless, we want to place an upper bound on the non-traded consumption risk premium in our model. In particular, we relax our assumption that that traded assets span all aggregate shocks by freeing up the 9^{th} element of Λ_0, the risk price of the non-traded consumption growth shock that is orthogonal to the eight traded asset shocks. Table 2 reports the consumption risk premium (Column 2), the average wealth-consumption ratio (Column 3), the maximum conditional Sharpe ratio (Column 4) and the Sharpe ratio on a one-period ahead consumption strip (Column 5) for different values of the price of non-traded consumption risk, governed by the 9th-element of Λ_0 (Column 1). This parameter does not affect the prices of any traded assets, so this exercise does not change any of the model's implications for observables.[19]

The first line reports our benchmark case in which the non-traded consumption risk is not priced. The consumption risk premium is 2.2% per annum, the maximum Sharpe ratio is 0.7 and the conditional Sharpe ratio on the one-period ahead consumption strip is .09. Increasing $\Lambda_0(9)$ increases the consumption risk premium, lowers the wealth-consumption ratio, and increases the Sharpe ratio on the consumption claim. How far should we increase $\Lambda_0(9)$? A first answer is to bound the maximal Sharpe ratio ($std_t[m_{t+1}]$). Cochrane and Saa-Requejo (2000) and Alvarez and Jermann (2004) choose a "good deal" bound of one, which they argue is high because it is twice the 0.5 Sharpe ratio on equities in the data.[20] Since we work with quarterly log returns, the

[19]Freeing up $\Lambda_0(9)$ also affects the risk premium and price-dividend ratio on human wealth, in quantitatively similar ways. We also experimented with freeing up the price of risk on the shock to labor income growth that is orthogonal to all previous shocks, including the aggregate consumption growth shock. Increasing this $\Lambda_0(10)$ has no effect on the consumption risk premium and the wealth-consumption ratio. It only affect the risk premium on human wealth. Quantitatively, those effects are similar to those presented in Table 2. The same is true when we simultaneously increase $\Lambda_0(9)$ and $\Lambda_0(10)$.

[20]In related work, Bernardo and Ledoit (2000) bound the gain-loss ratio which summarizes the attractiveness of a zero-price portfolio. It is equivalent to a restriction on admissible pricing kernels, precluding the existence of

Sharpe ratio on equities is only 0.22, and that same good deal bound of one is more than twice as conservative. This bound is reached for $\Lambda_0(9)$ around 0.8, and implies a consumption risk premium of 3.58% per annum and an average wealth-consumption ratio of 38. Even then, the consumption risk premium is still 4.3% short of the equity premium, so that our conclusion that total wealth has different risk-return characteristics than equity remains valid. In order to match the equity premium by increasing the price of non-traded consumption risk, we would need an increase in the maximum Sharpe ratio to three times the good-deal bound or 14 times the Sharpe ratio on equity.

A second answer would be to evaluate the Sharpe ratios on the consumption strip return in Column 5. When we set $\Lambda_0(9)$ to 0.1, this Sharpe ratio doubles compared to $\Lambda_0(9) = 0$. I.e., the implied price of non-traded consumption cash flow risk is much higher than that of traded consumption cash flow risk on a *per unit of risk* basis. Allowing the consumption strip to have the same magnitude Sharpe ratio as equity (0.22), would imply a value for $\Lambda_0(9)$ around 0.10. At this value the consumption risk premium is only about 0.2% per year higher than in our benchmark case. At $\Lambda_0(9) = 0.8$, the conditional Sharpe ratio on the consumption strip is 0.81, four times higher than the Sharpe ratio on equity and eight times higher than the Sharpe ratio on the traded consumption strip.

A third answer comes from utility-based pricing (Henderson 2002). We imagine that all traded risks are priced with the SDF we have estimated from stocks and bonds, but the non-traded consumption risk is evaluated with a constant relative risk aversion utility function (as usual in the utility-based pricing approach). Then the price of non-traded consumption risk equals $\Lambda_0(9) = \alpha\sigma_U$, where α is the risk aversion parameter and σ_U is the standard deviation of the unspanned consumption growth innovation.[21] Using $\Lambda_0(9)$ from the first column and the σ_U from the data, the last column of Table 2 reports the implied risk aversion parameter α. This risk aversion parameter increases steeply in $\Lambda_0(9)$. At $\Lambda_0(9) = .1$ it is 28 and at $\Lambda_0(9) = .8$ it is 220. For a risk aversion coefficient of 10, considered by many economists to be an upper bound, the consumption risk premium is 2.23%, a mere 4 basis points above our benchmark estimate. The same exercise with external habit instead of CRRA preferences leads to the same conclusion, but has the advantage that it can generate realistic equity premia. It implies a price of non-traded consumption risk of $\Lambda_0(9) = \alpha\bar{S}^{-1}\sigma_U$. In the benchmark external habit calibration (see Section 5.2), α is 2 and $\bar{S}^{-1}$ is 21, which implies that $\Lambda_0(9) = .15$. The corresponding consumption risk premium is 2.43%, only 24 basis points above our estimate.

arbitrage and approximate arbitrage opportunities.

[21]Formally, let the SDF be the product of two components $M_{t+1} = M_{t+1}^T M_{t+1}^{NT}$. The first component is the SDF that prices the traded asset from before: $\log M_{t+1}^T = y_0(1) - .5\Lambda_t(1:8)'\Lambda_t(1:8) - \Lambda_t(1:8)'\varepsilon_{t+1}(1:8)$. The second component is the SDF of a representative agent with CRRA utility which is orthogonalized on the traded SDF M_{t+1}^T and rescaled so that $E_t[M_{t+1}^{NT}] = 1$. The rescaling ensures that the risk-free rate continues to equal $y_0(1)$. Let the mean-zero innovation to consumption growth that is orthogonal to the first eight VAR shocks be $\sigma_U\varepsilon_{t+1}^U$, with $Var[\varepsilon_{t+1}^U] = 1$. Then $\log M_{t+1}^{NT} = -.5\alpha^2\sigma_U^2 - \alpha\sigma_U\varepsilon_{t+1}^U$. Equivalently, this imposes that the CRRA SDF prices that return on total wealth which is orthogonal to all traded asset returns.

[Table 2 about here.]

4.2 State Dynamics

The results of our estimation exercise are robust to different specifications of the law of motion for the state z. We consider three alternative models. Table 3 summarizes the key statistics for each of the specifications; the first row is the benchmark from the preceding analysis. In a first robustness exercise, labeled "simple return," we simplify the stock market dynamics. In particular, we assume that the log price-dividend ratio on equity pd^m follows an AR(1), that the expected aggregate stock return is only predicted by pd^m_t, that the fmp return for consumption is only predicted by pd^m_t and its own lag, and that the fmp return for labor income is only predicted by pd^m_t, the lagged fmp return for consumption, and its own lag. This zeroes out Ψ_{21} and Ψ_{31} and it simplifies the blocks Ψ_{22} and Ψ_{32} in the companion matrix. Because of the non-zero correlation between the shocks to the term structure and to the stock market variables, the prices of stock market risk inherit an exposure to the term structure variables, so that the elements of $\Lambda_{1,21}$ remain non-zero. The "simple return" specification shows very similar unconditional and conditional moments for the wc ratio. The last column shows a similar fit with the benchmark model; the sum of squared deviations between the moments in the model and in the data is 684 versus 676 in the benchmark. In a second robustness exercise, labeled "c-l predicts stocks", we replace log labor income growth Δl by the log consumption to labor income ratio $c - l$. This enables us to impose cointegration between the consumption and labor income streams. Just like $E_t[\Delta l_{t+1}]$ before, we assume that $E_t[c_{t+1} - l_{t+1}]$ depends on all VAR elements. In addition, lagged $c - l$ is also allowed to predict future consumption growth so that Ψ_{43} has non-zero elements everywhere. We keep the simplified structure for Ψ_{21}, Ψ_{22}, Ψ_{31}, and Ψ_{32} from the previous exercise, but we allow Δc and $c - l$ to predict future stock and fmp returns. That is, we free up Ψ_{34} and the last two elements in Ψ_{24}. Consumption growth and to a lesser extent the consumption-labor income ratio have significant predictive power for stock returns and the R^2 of the aggregate return equation increases from 7.6% (benchmark) to 10.6%. This predictability has also been found by Santos and Veronesi (2006) and Lettau and Ludvigson (2001a). Because of the change in Ψ_{24} and Ψ_{34}, this specification requires two (four) additional non-zero elements in $\Lambda_{1,24}$ ($\Lambda_{1,34}$). The third row of Table 3 shows that the wealth-consumption ratio properties are again similar. The mean wealth-consumption ratio is slightly higher and the total wealth return slightly more volatile. The extra flexibility improves the fit. The last exercise, labeled "c-l predicts yield" keeps the structure of the previous robustness exercise, but allows lagged aggregate consumption growth and the lagged consumption-labor ratio to predict the four term structure variables. This frees up Ψ_{14} and identifies four elements in $\Lambda_{1,14}$ ($\Lambda_1[2, 9 : 10]$ and $\Lambda_1[4, 9 : 10]$). The motivation is that a measure of real economic activity, such as consumption growth, is often included as a term structure determinant in the no-arbitrage term structure literature. The wealth-consumption ratio increases a bit further, but not the consumption risk premium. The reason is that the real yield curve is slightly less steep. In conclusion, the various specifications for Ψ and Λ_1 we explored lead to quantitatively similar results. The average consumption risk premium is in a narrow band between 2.16 and 2.24 percent per year; the same is true for the mean wealth-consumption ratio. All calibrations suggest mild predictability of total wealth returns. Whatever predictability there is comes from return predictability, not cash flow predictability. Finally, the future return predictability comes mostly from future risk-free rate predictability, except for the last calibration where risk-free rate predictability is somewhat less pronounced.

[Table 3 about here.]

4.3 Annual Estimation

Annual VAR dynamics may capture lower-frequency correlations between consumption growth and traded asset prices. To investigate this possibility, we have re-estimated the wealth-consumption ratio on annual data over the period 1952-2007.[22] We find that annual consumption growth has a significantly positively covariance with stock returns (t-stat is 3), which contributes to a better spanning of annual consumption growth risk by the traded assets than in the quarterly model. Nevertheless, the results from the annual estimation are similar to those of the quarterly model. The consumption risk premium is 1.7% (versus 2.2%) with a volatility of 22% (versus 17%). The dynamics of the wealth-consumption ratio still mirror those of long-term real bond yields. The variance decomposition still attributes all the variance of the wealth-consumption ratio to covariation with future total wealth returns rather than with future consumption growth. There is still some evidence that the wealth-consumption ratio predicts future risk-free rates, albeit weaker than in the quarterly model.

5 The Wealth-Consumption Ratio in Leading DAPMs

In the last part of the paper, we repeat the measurement exercise inside the two leading DAPMs: the long-run risk (LRR) and the external habit (EH) model. Just like in the model we estimate, the log wealth-consumption ratio is linear in the state variables in each of the models. We do *not* attempt to formally test the two models, only to point out their implications for the

[22] The only two differences with the quarterly model is that the short-term rate is now the 1-year Treasury bond yield and that we use the price-dividend ratio on the stock market without repurchase adjustment. The log pd^m ratio has a volatility of 40%. The CP factor is estimated on forward rates constructed form annual yields, and the factor mimicking portfolio returns are also constructed based on annual portfolio returns.

wealth-consumption ratio.[23] Interestingly, they have quite different implications for the wealth-consumption ratio. This difference is what motivated us to the measurement exercise above.

5.1 The Long-Run Risk Model

The long-run risk literature works off the class of preferences due to Kreps and Porteus (1978), Epstein and Zin (1989), and Duffie and Epstein (1992); see Appendix D.1. These preferences impute a concern for the timing of the resolution of uncertainty. A first parameter α governs risk aversion and a second parameter ρ governs the willingness to substitute consumption inter-temporally. In particular, ρ is the inverse of the inter-temporal elasticity of substitution (EIS). We adopt the consumption growth specification of Bansal and Yaron (2004):

$$\Delta c_{t+1} = \mu_c + x_t + \sigma_t \eta_{t+1}, \tag{14}$$

$$x_{t+1} = \rho_x x_t + \varphi_e \sigma_t e_{t+1}, \tag{15}$$

$$\sigma^2_{t+1} = \sigma^2 + \nu_1(\sigma^2_t - \sigma^2) + \sigma_w w_{t+1}, \tag{16}$$

where (η_t, e_t, w_t) are i.i.d. standard normal innovations. Consumption growth contains a low-frequency component x_t and is heteroscedastic, with conditional variance σ^2_t. The two state variables x_t and σ^2_t capture time-varying growth rates and time-varying economic uncertainty.

Proposition 3. *The log wealth-consumption ratio is linear in the two state variables* $z^{LRR}_t = [x_t, \sigma^2_t - \bar{\sigma}^2]$*:*

$$wc_t = A^{c,LRR}_0 + A^{c,LRR\prime}_1 z^{LRR}_t. \tag{17}$$

Appendix D.2 proves the proposition, following Bansal and Yaron (2004), and spells out the dependence of $A^{c,LRR}_0$ and $A^{c,LRR}_1$ on the structural parameters. This proposition implies that the log SDF in the LRR model can be written as a linear function of the growth rate of consumption and the growth rate of the log wealth-consumption ratio.[24] This two-factor representation highlights the importance of understanding the wc ratio dynamics for the LRR model's asset pricing implications.

We calibrate and simulate the long-run risk model choosing the benchmark parameter values of Bansal and Yaron (2004).[25] Column 1 of Table 1 reports the moments for the LRR model. All

[23]The LRR and EH models are not nested by our model. Their state displays heteroscedasticity, which translates into market prices of risk Λ_t are affine in the square root of the state. Our model has conditionally homoscedastic state dynamics and linear market prices of risk, but more shocks and therefore richer market price of risk dynamics.

[24]This result is formally stated and proven in Appendix D.2. Furthermore, appendix D.1 proves that the ability to write the SDF in the LRR model as a (non-linear) function of consumption growth and the wc ratio is general. It does not depend on the linearization of returns, nor on the consumption growth process in (14)-(16).

[25]Since their model is calibrated at monthly frequency but the data are quarterly, we work with a quarterly calibration instead. Appendix D.3 describes the mapping from monthly to quarterly parameters, the actual parameter values, and details on the simulation.

reported moments are averages across 5,000 simulations. The standard deviation of these statistics across simulations are bootstrap standard errors, and are reported in parentheses. The LRR model produces a wc ratio that is very smooth. Its volatility is 2.35%, quite a bit lower than in the data (last column). Almost all the volatility in the wealth-consumption ratio comes from volatility in the persistent component of consumption (the volatility of x is about 0.5% and the loading of wc on x is about 5). The persistence of both state variables induces substantial persistence in the wc ratio: its auto-correlation coefficient is 0.91 (0.70) at the 1-quarter (4-quarter) horizon. The change in the wc ratio, which is the second asset pricing factor in the log SDF, has a volatility of 0.90%. Aggregate consumption growth, the first asset pricing factor, has a higher volatility of 1.45%. The correlation between the two asset pricing factors is statistically indistinguishable from zero. The resulting log total wealth return has a volatility of 1.64% per quarter in the LRR model, again lower than in the data. Low autocorrelation in Δwc and Δc generates low autocorrelation in total wealth returns. The total wealth return has a counter-factually high correlation with consumption growth (+.84) because most of the action in the total wealth return comes from consumption growth. The lower panel reports the consumption risk premium, the expected return on total wealth in excess of the risk-free rate (including a Jensen term). Total wealth is not very risky in the LRR model; the quarterly risk premium is 40 basis points, which translates into 1.6% per year. Each asset pricing factor contributes about half of the risk premium. The low consumption risk premium corresponds to a high average wealth-consumption ratio; it is 87 expressed in annual levels ($e^{A_0^{c,LRR}-\log(4)}$). Just as in the data, total wealth is not very risky in the LRR model.

Turning to the conditional moments, the amount of total wealth return predictability is low because the wealth-consumption ratio is smooth. The (demeaned) wc ratio can be decomposed into a discount rate and a cash flow component:

$$wc_t = \Delta c_t^H + r_t^H = \left[\frac{1}{\kappa_1^c - \rho_x} x_t\right] - \left[\frac{\rho}{\kappa_1^c - \rho_x} x_t - A_2^{LRR}\left(\sigma_t^2 - \overline{\sigma}^2\right)\right].$$

Appendix D.4 derives this decomposition as well as the decomposition of the variance of wc. The discount rate component itself contains a risk-free rate component and a risk premium component. The persistent component of consumption growth x_t drives only the risk-free rate effect (first term in r_t^H). It is governed by ρ, the inverse EIS. In the log case ($\rho = 1$), the cash flow loading on x and the risk-free rate loading on x exactly offset each other. The risk premium component is driven by the heteroscedastic component of consumption growth.[26] The expressions for the theoretical covariances of wc_t with Δc_t^H and $-r_t^H$ show that both cannot simultaneously be positive. When $\rho < 1$, the sign on the regression coefficient of future consumption growth on the log wealth-consumption ratio is positive, but the sign on the return predictability equation is negative (unless

[26]The heteroscedasticity also affects the risk-free rate component, but without heteroscedasticity there would be no time-variation in risk premia.

the heteroscedasticity mechanism is very strong). The opposite is true for $\rho > 1$ (low EIS). The benchmark calibration of the LRR model has a high EIS. Most of the volatility in the wealth-consumption ratio arises from covariation with future consumption growth (297.5%). The other -197.5% is accounted for by the covariance with future returns. A calibration with an EIS below 1 would generate the same sign on the covariance with returns as in the data. Alternatively, a positive correlation between innovations to x and $\sigma_t^2 - \bar{\sigma}^2$ may help to generate a variance decomposition closer to the data. Finally, virtually all predictability in future total wealth returns arises from predictability in future risk-free rates. This is similar to what we find in the data.

Despite the low consumption risk premium and high wc ratio, the LRR model is able to match the high equity risk premium and low pd^m ratio. The reason is that the dividend claim carries more long run risk: dividend growth has a loading of 3 on x_t whereas consumption growth only has a loading of 1.[27] Therefore, the LRR model generates the wedge between total wealth and equity we also find in the data.

5.2 The External Habit Model

We use the specification of preferences proposed by Campbell and Cochrane (1999), henceforth CC. The log SDF is

$$m_{t+1} = \log \beta - \alpha \Delta c_{t+1} - \alpha (s_{t+1} - s_t),$$

where the log surplus-consumption ratio $s_t = \log(S_t) = \log\left(\frac{C_t - X_t}{C_t}\right)$ measures the deviation of consumption C_t from the habit X_t, and has the following law of motion:

$$s_{t+1} - \bar{s} = \rho_s (s_t - \bar{s}) + \lambda_t (\Delta c_{t+1} - \mu_c).$$

The steady-state log surplus-consumption ratio is $\bar{s} = \log\left(\bar{S}\right)$. The parameter α continues to capture risk aversion. The "sensitivity" function λ_t governs the conditional covariance between consumption innovations and the surplus-consumption ratio and is defined below in (20). As in CC, we assume an i.i.d. consumption growth process:

$$\Delta c_{t+1} = \mu_c + \bar{\sigma} \eta_{t+1}, \tag{18}$$

where η is an i.i.d. standard normal innovation and the only shock in the model.

Just as in the LRR model and in the data, the log wealth-consumption ratio is affine in the

[27]See Appendix D.5 for the dividend growth specification, and the expressions for the log price-dividend ratio on equity, and the equity risk premium. The Bansal and Yaron (2004) calibration of dividend growth does not impose cointegration between consumption and dividends. Bansal, Dittmar, and Lundblad (2005), Bekaert, Engstrom, and Grenadier (2005) and Bekaert, Engstrom, and Xing (2008) consider versions of the LRR model with cointegration. The results with cointegration are similar and are omitted for brevity.

state variable of the EH model.

Proposition 4. *The log wealth-consumption ratio is linear in the sole state variable* $z_t^{EH} = s_t - \bar{s}$*,*

$$wc_t = A_0^{c,EH} + A_1^{c,EH} z_t^{EH}, \tag{19}$$

and the sensitivity function takes the following form

$$\lambda_t = \frac{\bar{S}^{-1}\sqrt{1-2(s_t-\bar{s})}+1-\alpha}{\alpha - A_1} \tag{20}$$

Appendix E.1 proves this proposition. Just like CC's sensitivity function delivers a risk-free rate that is linear in the state $s_t - \bar{s}$, our sensitivity function delivers a log wealth-consumption ratio that is linear in $s_t - \bar{s}$. To minimize the deviations with the CC model, we pin down the steady-state surplus-consumption level $\bar{S}$ by matching the steady-state risk-free rate to the one in the CC model. Taken together with the expressions for $A_0^{c,EH}$ and $A_1^{c,EH}$, this restriction amounts to a system of three equations in three unknowns $\left(A_0^{c,EH}, A_1^{c,EH}, \bar{S}\right)$.[28] This proposition implies that the log SDF in the EH model is a linear function of the same two asset pricing factors as in the LRR model: the growth rate of consumption and the growth rate of the consumption-wealth ratio. Appendix E.1 shows this result more formally. Therefore, differences in the properties of the wc ratio between models generates differences in their asset pricing predictions.

We calibrate the EH model choosing the benchmark parameter values of CC.[29] The simulation method is parallel to the one described for the LRR model. We note that the risk-free rate is nearly constant in the benchmark calibration; its volatility is .03% per quarter. This shows that the slight modification in the sensitivity function from the CC one did not materially alter the properties of the risk-free rate. The second column of Table 1 reports the moments of the wealth-consumption ratio under the benchmark calibration of the EH model. First and foremost, the wc ratio is volatile in the EH model: it has a standard deviation of 29.3%, which is 12.5 times larger than in the LRR model and 12 percentage points higher than in the data. This volatility comes from the high volatility of the surplus consumption ratio (38%). The persistence in the surplus-consumption ratio drives the persistence in the wealth-consumption ratio: its auto-correlation coefficient is 0.93 (0.74) at the 1-quarter (4-quarter) horizon. The change in the wc ratio has a volatility of 9.46%. This is more than 10 times higher than the volatility of the first asset pricing factor, consumption growth, which has a standard deviation of 0.75%. The high volatility of the change in the wc ratio translates into a highly volatile total wealth return. The log total wealth return has a volatility of 10.26% per quarter in the EH model. As in the LRR model, the total wealth return is strongly positively correlated with consumption growth (.91). In the EH model this happens because most

[28] Details are in Appendix E.2. Appendix E.3 discusses an alternative way to pin down $\bar{S}$.

[29] Appendix E.4 describes the mapping from monthly to quarterly parameters and reports the parameter values.

of the action in the total wealth return comes from changes in the wc ratio. The latter are highly positively correlated with consumption growth (.90, in contrast with the LRR model). Finally, the consumption risk premium is high because total wealth is risky; the quarterly risk premium is 267 basis points, which translates into 10.7% per year. Most of the risk compensation in the EH model is for bearing Δwc risk. The high consumption risk premium implies a low mean log wealth-consumption ratio of 3.86. Expressed in annual levels, the mean wealth-consumption ratio is 12.

In contrast to the LRR model, the EH model asserts that all variability in returns arises from variability in risk premia (see Appendix E.5). Since there is no consumption growth predictability, 100% of the variability of wc is variability of the discount rate component. The covariance between the wealth-consumption ratio and returns has the right sign: it is positive by construction. This variance decomposition is close to the data. A key strength of the EH model is its ability to generate a lot of variability in expected equity returns. The flip side is that the same mechanism also generates a lot of variability in expected total wealth returns. Finally, the EH model implies that almost all the covariance with future returns comes from covariance with future excess returns, not future risk-free rates. In the data, there was evidence for risk-free rate predictability.

The properties of total wealth returns are similar to those of equity returns.[30] The equity risk premium is only 1.2 times higher than the consumption risk premium and the volatility of the pd^m ratio is only 1.2 times higher than the volatility of the wc ratio. For comparison, in the LRR model, these ratios are 3.5 and 6 and in the data they are 3.3 and 1.6, respectively. The EH model drives not enough of a wedge between the riskiness of total wealth and equity.

In sum, the two leading asset pricing models have very different implications for the wealth consumption ratio, despite the fact that they both match unconditional equity return moments.[31] In the LRR model, the consumption claim looks more like a bond whereas in the EH model it looks more like a stock.

6 Conclusion

We develop a new methodology for estimating the wealth-consumption ratio in the data, based on no-arbitrage conditions that are familiar from the term structure literature. It combines restrictions on stocks and bonds in a novel way. We find that a claim to aggregate consumption is much less risky than a claim to aggregate dividends: the consumption risk premium is only one-third of the equity risk premium. This suggests that the stand-in households' portfolio is much less risky than

[30] Appendix E.6 contains the details of the dividend growth specification, the calibration, and the computation of the price-dividend ratio and equity returns. The dividend growth specifications in Campbell and Cochrane (1999) and Wachter (2006) do not impose cointegration with consumption growth, while Wachter (2005) does.

[31] A previous version of this paper contained simulated method of moments estimation of both models, targeting both total wealth and equity return moments.

what one would conclude from studying the equity component of that portfolio. The consumption claim looks much more like a real bond than like a stock.

Our findings have clear implications for future work on dynamic asset pricing models. In any model, the same stochastic discount factor needs to price both a claim to aggregate consumption, which is not that risky and carries a low return, and a claim to equity dividends, which is much more risky and carries a high return. Generating substantial time-variation in expected equity returns though variation in conditional market prices of risk may have the effect of generating too much time-variation in expected total wealth returns. Our exercise suggests that stocks are special, so that predictability in equity returns may need to be generated through an interesting correlation structure between the cash flow process and the stochastic discount factor.

References

Abel, A. B. (1990): "Asset Prices Under Habit Formation and Catching Up With the Jones," *American Economic Review*, 80, 38–42.

——— (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33.

Alvarez, F., and U. Jermann (2004): "Using Asset Prices to Measure the Cost of Business Cycles.," *Journal of Political Economy*, 112(6), 1223–1255.

Ang, A., G. Bekaert, and M. Wei (2008): "The Term Structure of Real Rates and Expected Inflation," *Journal of Finance*, 63 (2), 797–849.

Ang, A., and J. Liu (2007): "Risk, Return, and Dividends," *Journal of Financial Economics*, 85, 1–38.

Ang, A., and M. Piazzesi (2003): "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics*, 50, 745–787.

Bansal, R., R. F. Dittmar, and C. Lundblad (2005): "Consumption, Dividends and the Cross-Section of Equity Returns," *Journal of Finance*, 60(4), 1639 – 1672.

Bansal, R., R. Gallant, and G. Tauchen (2007): "Rational Pessimism, Rational Exuberance, and Asset Pricing Models," *Review of Economic Studies*, forthcoming.

Bansal, R., D. Kiku, and A. Yaron (2006): "Risks for the Long Run: Estimation and Inference," Working Paper Duke University.

Bansal, R., and I. Shaliastovich (2007): "Risk and Return in Bond, Currency, and Equity Markets," Working Paper Duke University.

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

Bekaert, G., E. Engstrom, and S. Grenadier (2005): "Stock and Bond Returns with Moody Investors," Working Paper Columbia University.

Bekaert, G., E. Engstrom, and S. Xing (2008): "Risk, Uncertainty and Asset Prices," Forthcoming Journal of Financial Economics.

Bekaert, G., and S. R. Grenadier (2001): "Stock and Bond Pricing in an Affine Economy," NBER Working Paper No. 7346.

Benzoni, L., P. Collin-Dufresne, and R. S. Goldstein (2008): "Explaining Pre- and Post-1987 Crash Prices of Equity and Options within a Unified General Equilibrium Framework," Working Paper University of Minnesota.

Bernardo, A. E., and O. Ledoit (2000): "Gain, Loss and Asset Pricing," *Journal of Political Economy*, 108 (1), 144–172.

Binsbergen, J., and R. Koijen (2007): "Predictive Regressions: A Present-Value Approach," Working Paper University of Chicago and Stanford University.

Boudoukh, J., R. Michaely, M. Richardson, and M. Roberts (2004): "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing," *NBER Working Paper*, 10651.

Breeden, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Campbell, J., A. Sunderam, and L. Viceira (2007): "Inflation Bets or Deflation Hedges? The Changing Risk of Nominal Bonds," Working Paper Harvard University.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (1993): "Intertemporal Asset Pricing Without Consumption Data," *American Economic Review*, 83(3), 487–511.

——— (1996): "Understanding Risk and Return," *The Journal of Political Economy*, 104(2), 298–345.

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

Campbell, J. Y., and R. J. Shiller (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–227.

Chen, L., P. Collin-Dufresne, and R. S. Goldstein (2008): "On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle," *Review of Financial Studies*, Forthcoming.

Chen, X., J. Favilukis, and S. Ludvigson (2008): "An Estimation of Economic Models with Recursive Preferences," Working Paper New York University.

Chen, X., and S. C. Ludvigson (2007): "Land of Addicts? An Empirical Investigation of Habit-Based Asset Pricing Models," Working Paper New York University.

Chien, Y., H. Cole, and H. Lustig (2007): "Macro Implications of Household Finance," Working Paper UCLA.

Cochrane, J. H. (2001): *Asset Pricing*. Princeton University Press, Princeton, N.J.

Cochrane, J. H., and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95, 138–160.

Cochrane, J. H., and J. Saa-Requejo (2000): "Beyond Arbitrage: 'Good deal' asset price bounds in incomplete markets," *Journal of Political Economy*, 108, 79–119.

Colacito, R., and M. Croce (2005): "Risks for the Long-Run and the Real Exchange Rate," Working Paper New York University.

Constantinides, G. M. (1990): "Habit-formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, 519–543.

Constantinides, G. M., and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.

Cuoco, D., and H. He (2001): "Dynamic Equilibrium in Infinite Dimensional Economies with Incomplete Financial Markets," *Annals of Economics and Finance*, 2, 265–296, Annals of Economics and Finance.

Dai, Q., and K. J. Singleton (2000): "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943–1978.

Duffie, D., and L. G. Epstein (1992): "Stochastic Differential Utility," *Econometrica*, 60(2), 353–394.

Duffie, D., and R. Kan (1996): "A Yield Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.

Epstein, L., and S. Zin (1989): "Substitution Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–968.

Fama, E. F., and K. R. French (1992): "The Cross-Section of Expected Returns," *Journal of Finance*, 47, 427–465.

Ferson, W. E., and G. M. Constantinides (1991): "Habit Persistence and Durability in Aggregate Consumption," *Journal of Financial Economics*, 29, 199–240.

Ferson, W. E., and C. R. Harvey (1991): "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385–415.

Hansen, L. P., J. Heaton, J. Lee, and N. Roussanov (2008): "Intertemporal Substitution and Risk Aversion," in *Handbook of Econometrics*, ed. by J. Heckman, and E. Leamer, vol. 6, chap. 61. Elsevier Science.

Hansen, L. P., J. C. Heaton, and N. Li (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116 (2), 260–302.

Harvey, C. R. (1988): "The Real Term Structure and Consumption Growth," *Journal of Financial Economics*, 22, 305–333.

Henderson, V. (2002): "Valuation of Claims on Nontraded Asset Using Utility Maximization," *Mathematical Finance*, 12, 351–373.

Jagannathan, R., and Z. Wang (1996): "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance*, 51, 3–54.

Jorgenson, D. M., and B. Fraumeni (1989): "The Accumulation of Human and Non-Human Capital," in *The Measurement of Saving, Investment and Wealth*, ed. by R. Lipsey, and H. Tice, pp. 227–282. University of Chicago press, University of Chicago.

Kreps, D., and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Lettau, M., and S. C. Ludvigson (2001a): "Consumption, Aggregate Wealth and Expected Stock Returns," *The Journal of Finance*, 56(3), 815–849.

——— (2001b): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *The Journal of Political Economy*, 109(6), 1238–1287.

Lettau, M., and S. C. Ludvigson (2005): "Expected Returns and Expected Dividend Growth," *Journal of Financial Economics*, 76, 583–626.

Lettau, M., and S. Van Nieuwerburgh (2007): "Reconciling the Return Predictability Evidence," Forthcoming Review of Financial Studies.

Lettau, M., and J. Wachter (2007): "The Term Structures of Equity and Interest Rates," Working Paper University of Pennsylvania Wharton School.

Lucas, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–54.

Lustig, H., and Y.-L. Chien (2007): "The Market Price of Aggregate Risk and the Wealth Distribution," http://ssrn.com/abstract=290917.

Lustig, H., and S. Van Nieuwerburgh (2007a): "Can Housing Collateral Explain Long-Run Swings in Asset Returns?," Working Paper NYU Stern and UCLA.

——— (2007b): "The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street," *Forthcoming Review of Financial Studies.*

Malloy, C., T. Moskowitz, and A. Vissing-Jorgensen (2009): "Long-Run Stockholder Consumption Risk and Asset Returns," *Journal of Finance*, forthcoming.

Martin, I. (2008): "Consumption-based Asset Pricing with Higher-Order Cumulants," Working Paper Harvard University.

Menzly, L., T. Santos, and P. Veronesi (2004): "Understanding Predictability," *Journal of Political Economy*, 112(1), 1–47.

Parker, J., and C. Julliard (2005): "Consumption Risk and The Cross-Section of Expected Returns," *The Journal of Political Economy*, 113(1), 185–222.

Piazzesi, M., and M. Schneider (2006): "Equilibrium Yield Curves," *National Bureau of Economic Analysis Macroeconomics Annual.*

Roll, R. (1977): "A Critique of the Asset Pricing Theory's Tests: Part I," *Journal of Financial Economics*, 4, 129–176.

Santos, T., and P. Veronesi (2006): "Labor Income and Predictable Stock Returns," *Review of Financial Studies*, 19.

Shiller, R. J. (1995): "Aggregate Income Risks and Hedging Mechanisms," *Quarterly Review of Economics and Finance.*

Verdelhan, A. (2007): "A Habit-Based Explanation of the Exchange Rate Risk Premium," Working Paper Boston University.

Wachter, J. (2005): "Solving Models with External Habit," *Finance Research Letters*, 2, 210–226.

——— (2006): "A Consumption-Based Model of the Term Structure of Interest Rates," *Journal of Financial Economics*, 79, 365–399.

Yu, J. (2007): "The Long and the Short of Asset Prices: Using Long Run Consumption-Return Correlations to Test Asset Pricing Models," Working Paper Wharton School of Business.

Table 1: Moments of the Wealth-Consumption Ratio

This table displays unconditional moments of the log wealth-consumption ratio wc, its first difference Δwc, and the log total wealth return r^c. The last but one row reports the time-series average of the conditional consumption risk premium, $E[E_t[r_t^{c,e}]]$, where $r^{c,e}$ denotes the expected log return on total wealth in excess of the risk-free rate and corrected for a Jensen term. The first column reports moments from the long-run risk model (LRR model), simulated at quarterly frequency. All reported moments are averages and standard deviations (in parentheses) across the 5,000 simulations of 220 quarters of data. The second column reports the same moments for the external habit model (EH model). The last column is for the data. The standard errors are obtained by bootstrap, as described at the end of Appendix C.7.

Moments	LRR Model	EH model	data
$Std[wc]$	2.35%	29.33%	17.24%
(s.e.)	(.43)	(12.75)	(4.30)
$AC(1)[wc]$	.91	.93	.96
(s.e.)	(.03)	(.03)	(.03)
$AC(4)[wc]$	.70	.74	.85
(s.e.)	(.10)	(.11)	(.08)
$Std[\Delta wc]$	0.90%	9.46%	4.86%
(s.e.)	(.05)	(2.17)	(1.16)
$Std[\Delta c]$	1.43%	.75%	.44%
(s.e.)	(.08)	(.04)	(.03)
$Corr[\Delta c, \Delta wc]$	-.06	.90	.11
(s.e.)	(.06)	(.03)	(.06)
$Std[r^c]$	1.64%	10.26%	4.94%
(s.e.)	(.09)	(2.21)	(1.16)
$Corr[r^c, \Delta c]$	.84	.91	.19
(s.e.)	(.02)	(.03)	(.07)
$E[E_t[r_t^{c,e}]]$	0.40%	2.67%	0.54%
(s.e.)	(.01)	(1.16)	(.16)
$E[wc]$	5.85	3.86	5.86
(s.e.)	(.01)	(.17)	(.49)

Table 2: Non-traded Consumption Risk

The first column reports the market price of risk $\Lambda_0(9)$ that is associated with the innovation to consumption growth that is orthogonal to all innovations to the preceding stock and bond innovations. The second column reports the consumption risk premium. The third column reports the average wealth/consumption ratio. The fourth column is the maximum Sharpe ratio computed as $\sqrt{\Lambda_0'\Lambda_0}$. The last but one column shows the conditional Sharpe ratio on a one-period ahead consumption strip: $\left(e_c'\Sigma^{1/2}\Lambda_0\right)/\sqrt{e_c'\Sigma^{1/2}\Sigma^{1/2'}e_c}$. The last column reports the coefficient of relative risk aversion α for a CRRA investor so that the Euler equation for the consumption claim is satisfied.

$\Lambda_0(9)$	cons. risk premium	$E[WC]$	$std_t(m_{t+1})$	SR on strips	α
0	2.19%	87	0.69	0.09	0
0.05	2.26%	81	0.70	0.14	14
0.1	2.37%	75	0.70	0.18	28
0.5	3.05%	48	0.85	0.54	138
0.8	3.58%	38	**1.06**	0.81	220
1.0	3.94%	33	1.21	0.99	275
1.5	4.86%	25	1.65	1.44	413
2.0	5.79%	20	2.11	1.89	551
3.0	7.68%	15	3.08	2.79	826

Table 3: Robustness Analysis

The table reports the unconditional standard deviation of the log wealth-consumption ratio wc, the unconditional standard deviation of the log total wealth return r^c, the average consumption risk premium $E[E_t[r_t^{c,e}]]$ in percent per year, the mean log wealth-consumption ratio, the percentage of the variance of wc that is attributable to covariation of wc with future consumption growth ($pred_{CF} = Cov[wc_t, \Delta c_t^H]/Var[wc_t]$), and the percentage of the variance of wc that is attributable to covariation of wc with future risk-free rates $pred_rf$. The last column denotes the objection function value at the point estimate (obj).

Specifications	$Std[wc]$	$Std[r^c]$	$E[E_t[r_t^{c,e}]]$	$E[wc]$	$pred_{CF}$	$pred_{rf}$	obj
benchmark	17.24%	4.94%	0.54%	5.86	0.3%	69.6%	675.7
simple return	17.43%	4.89%	0.56%	5.81	0.4%	69.1%	684.3
c-l predicts stocks	18.00%	5.55%	0.55%	5.93	9.6%	61.7%	650.7
c-l predicts yield	19.10%	5.80%	0.56%	5.99	1.3%	22.6%	671.8

Figure 1: Average Term Structure of Interest Rates

The figure plots the observed and model-implied nominal bond yields for bonds of maturities 1-120 quarters. The data are obtained by using a spline-fitting function through the observed maturities. The third panel plots the model-implied real yields.

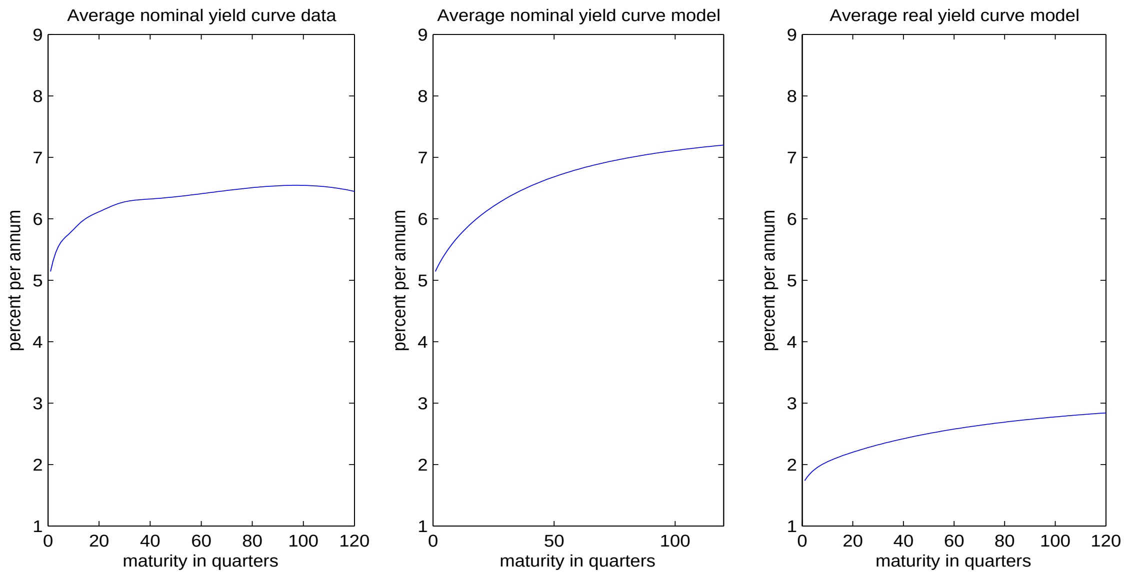

Figure 2: Dynamics of the Nominal Term Structure of Interest Rates

The figure plots the observed and model-implied 1-, 4-, 12-, 20-, 40-, and 80-quarter nominal bond yields. Note that the 20-year yield is unavailable between 1986.IV and 1993.II.

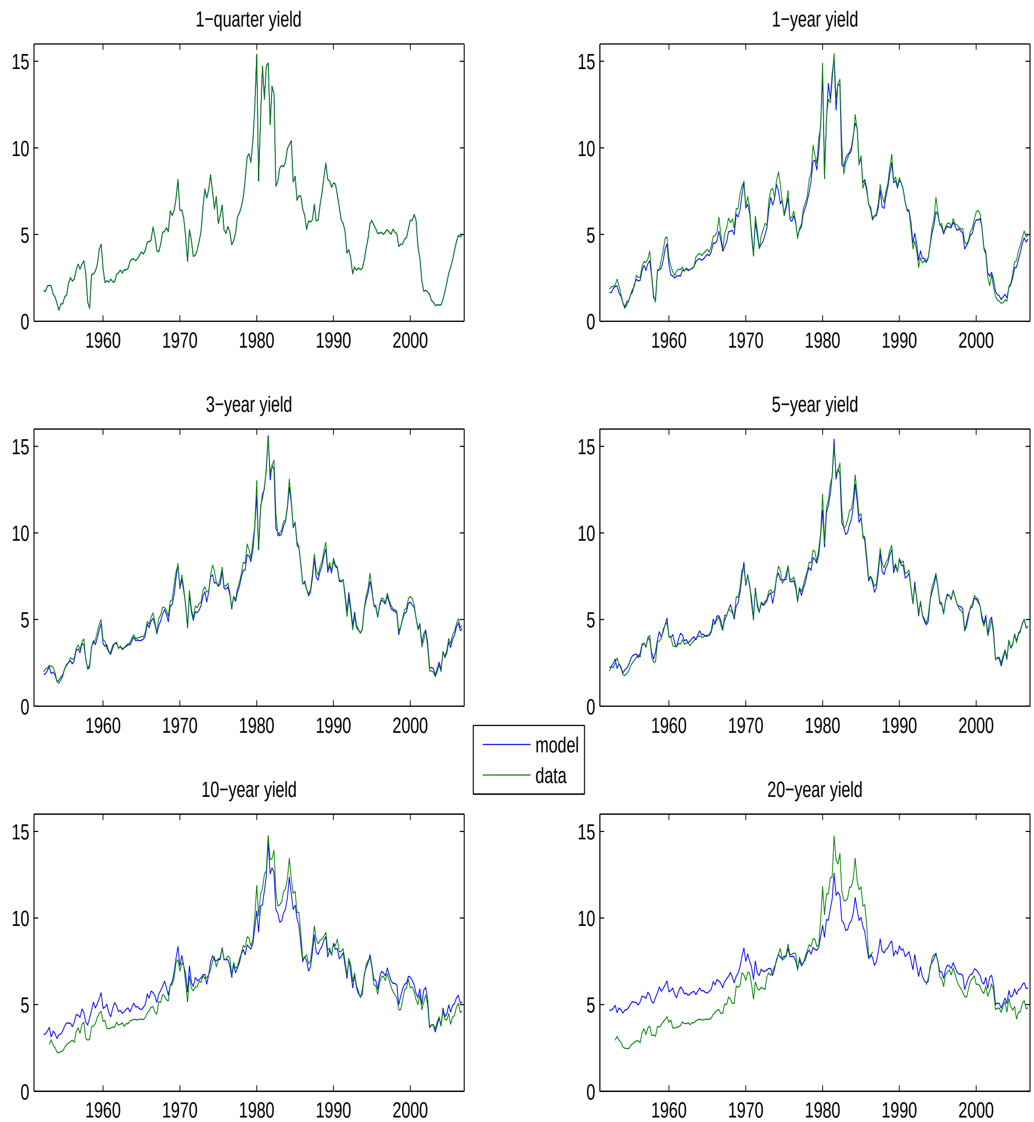

Figure 3: Nominal Bond Risk Premia

The left panel plots the 5-year nominal bond risk premium on a 5-year nominal bond in model and data. It is defined as the difference between the nominal 5-year yield and the expected future 1-quarter yield averaged over the next 5 years. It represents the return on a strategy that buys and holds a 5-year bond until maturity and finances this purchase by rolling over a 1-quarter bond for 5 years. The right panel plots the Cochrane-Piazzesi factor in model and data. It is a linear combination of the one-year nominal yield and 2- through 5-year forward rates. This linear combination is a good predictor of the one-quarter bond risk premium.

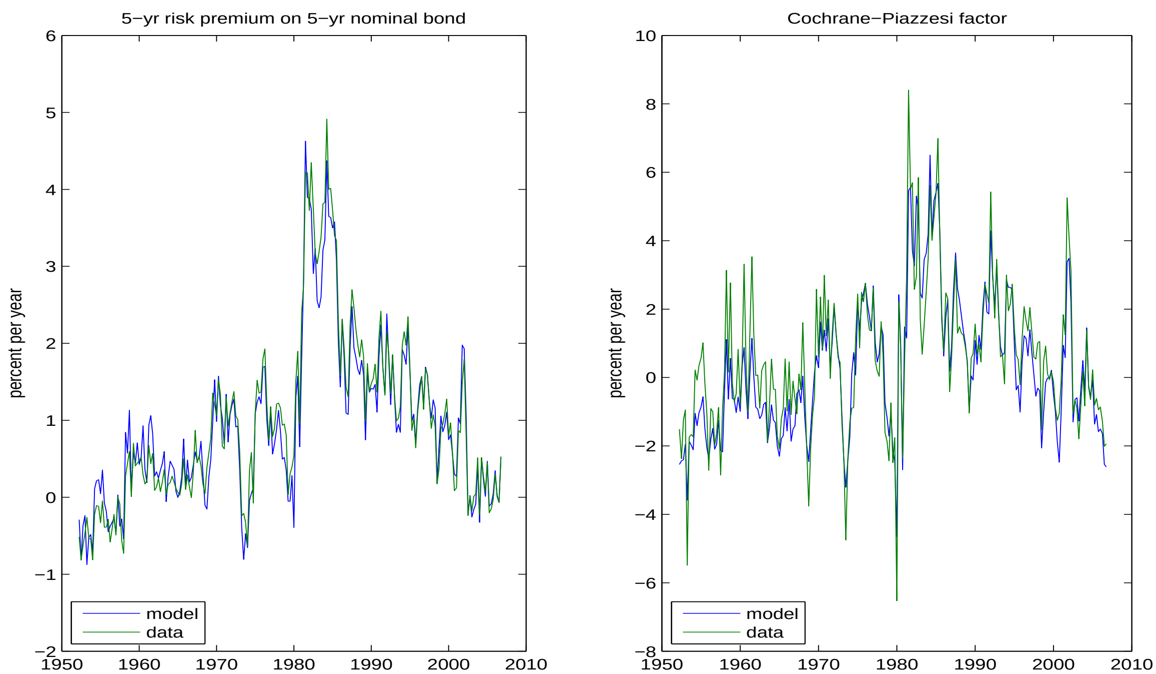

Figure 4: The Stock Market

The figure plots the observed and model-implied price-dividend ratio and expected excess return on the overall stock market.

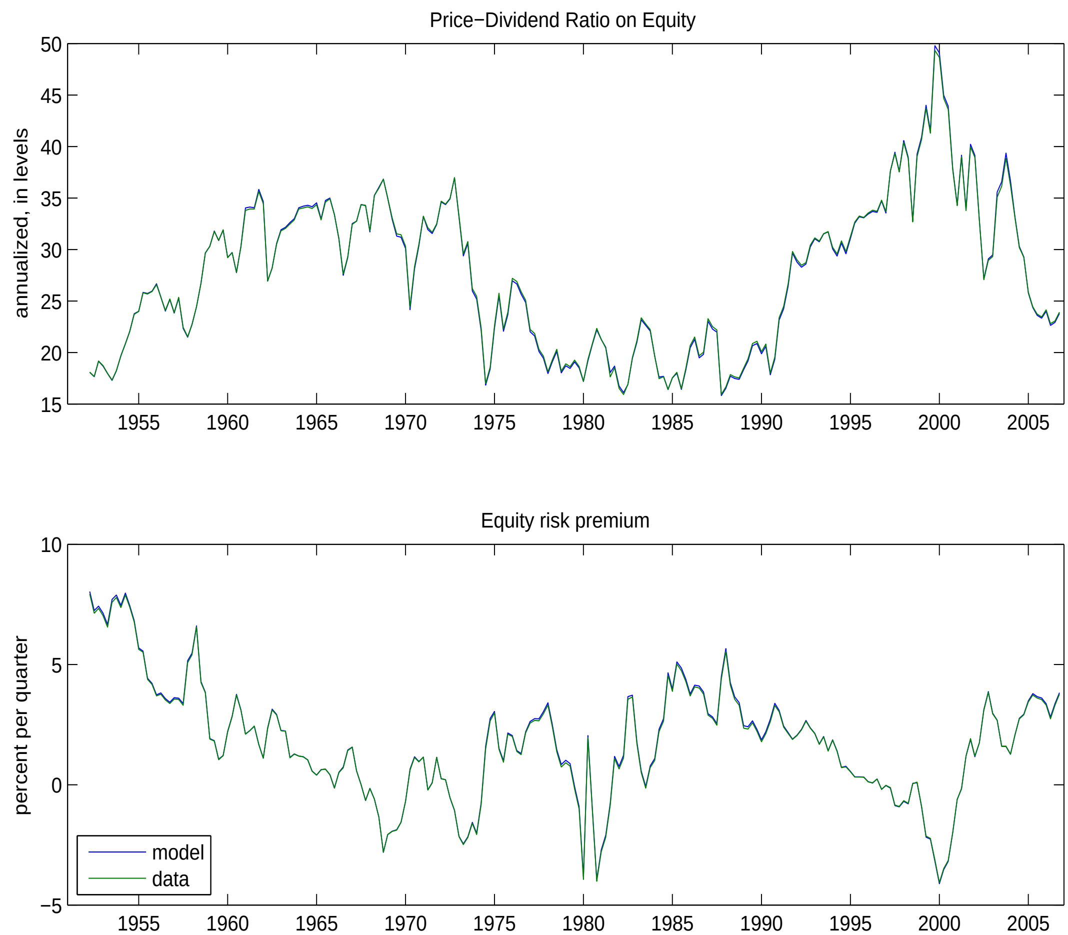

Figure 5: Decomposing the 5-Year Nominal Yield

The left panel decomposes the 5-year yield into the real 5-year yield, expected inflation over the next 5-years, and the inflation risk premium. The right panel decomposes the average nominal bond risk premium into the average real rate risk premium and inflation risk premium for maturities ranging from 1 to 120 quarters. The nominal (real) bond risk premium at maturity τ is defined as the nominal (real) τ-quarter yield minus the average expected future nominal (real) 1-quarter yield over the next τ quarters. The τ-quarter inflation risk premium, labeled as IRP, is the difference between the τ-quarter nominal and real risk premia.

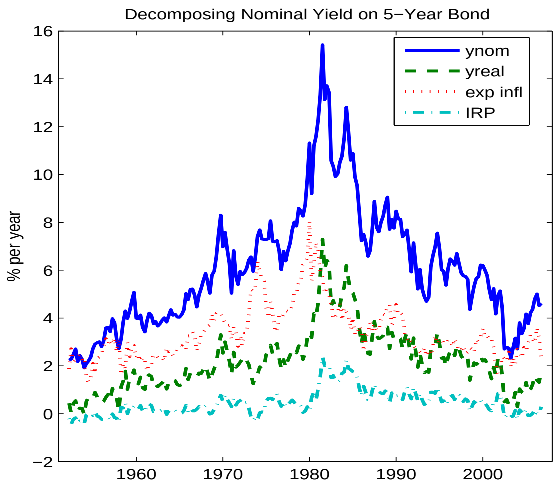

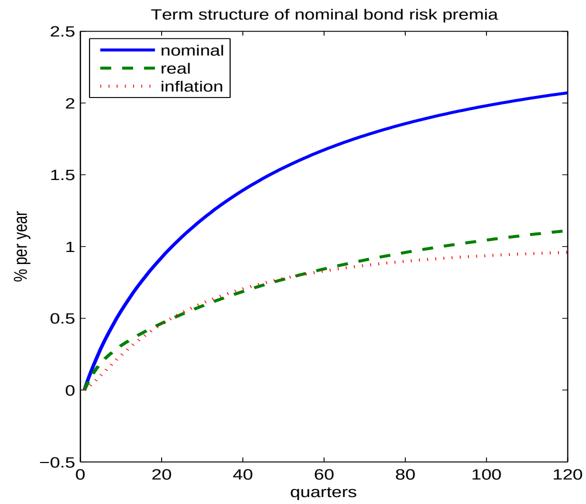

Figure 6: The Log Wealth-Consumption Ratio in the Data

The figure plots $\exp\{wc_t - log(4)\}$, where wc_t is the quarterly log total wealth to total consumption ratio. The log wealth consumption ratio is given by $wc_t = A_0^c + (A_1^c)' z_t$. The coefficients A_0^c and A_1^c satisfy equations (6)-(7).

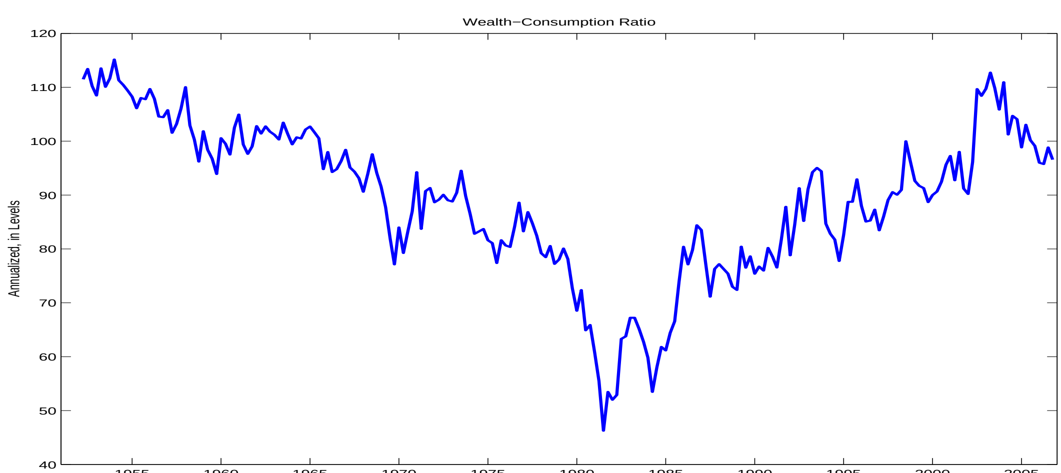

Figure 7: Discount Rates on Consumption and Dividend Claim

The figure plots the discount rate on a claim to consumption (solid line, measured against the left axis, in percent per year), the discount rate on a claim to dividend growth (dashed line, measured against the right axis, in percent per year), and the yield on a real 50-year bond (dotted line, measured against the right axis, in percent per year). The discount rates are the rates that make the price-dividend ratio equal to the expected present-discounted value of future cash flows, for either the consumption claim or the dividend claim.

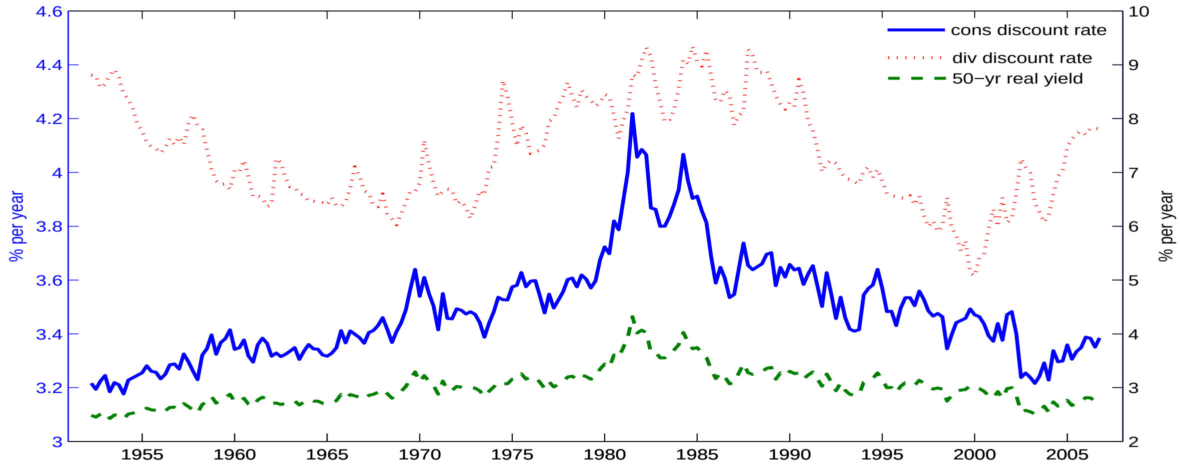

Figure 8: Decomposing the Yield on A Consumption Strip

The figure decomposes the yield on a consumption strip of maturity τ, which goes from 1 to 120 quarters, into a real bond yield minus deterministic consumption growth on the one hand and the yield on a security that only carries the consumption cash flow risk on the other hand. See C.5 for a detailed discussion of this decomposition.

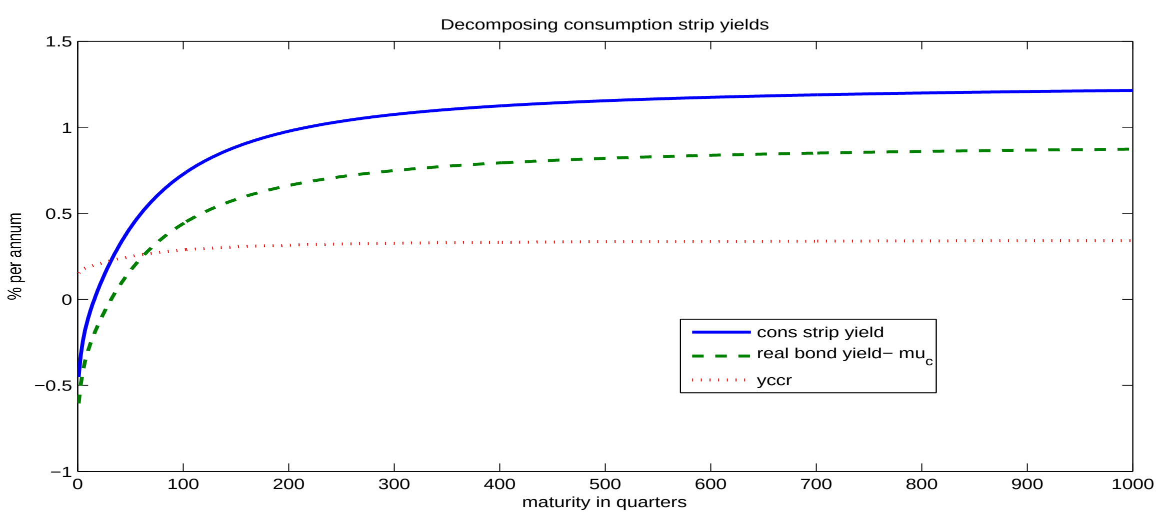

Figure 9: Real Per Capita Wealth Estimates

The left panel of the figure plots total wealth and human wealth as estimated from the data. The right panel plots their difference, which we label non-human wealth. It also plots the present discounted value of the first 35 years of labor income.

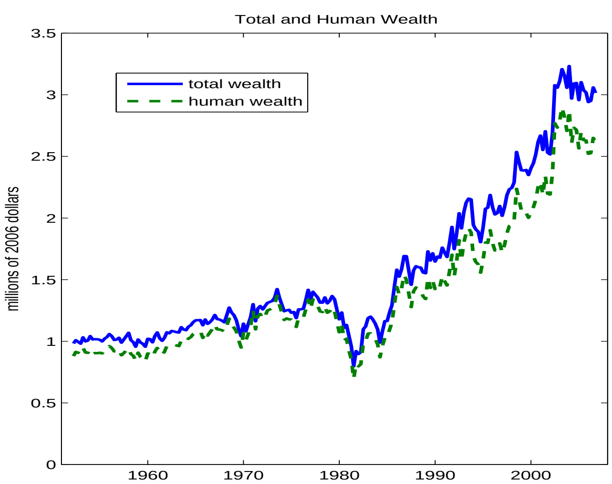

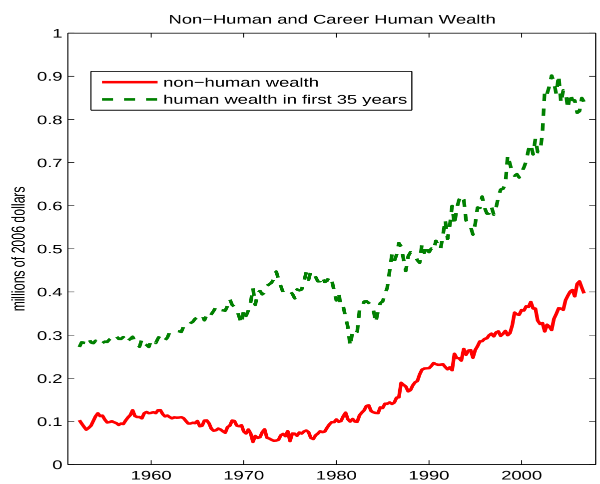

Technical Appendix for "The Wealth-Consumption Ratio"

A Pricing Aggregate Consumption and Labor Income

Environment Let $z_t \in Z$ be the aggregate state vector. We use z^t to denote the history of aggregate state realizations. Section 1.3 describes the dynamics of the aggregate state z_t of this economy, including the dynamics of aggregate consumption $C_t(z^t)$ and aggregate labor income $L_t(z^t)$.

We consider an economy that is populated by a continuum of heterogeneous agents, whose labor income is subject to idiosyncratic shocks. The idiosyncratic shocks are denoted by $\ell_t \in \mathcal{L}$, and we use ℓ^t to denote the history of these shocks. The household labor income process is given by:

$$\eta_t(\ell^t, z^t) = \widehat{\eta}_t(\ell^t, z_t) L_t(z^t).$$

Let $\Phi_t(z^t)$ denote the distribution of household histories ℓ^t conditional on being in aggregate node z^t. The labor income shares $\hat{\eta}$ aggregate to one:

$$\int \widehat{\eta}_t(\ell^t, z_t) d\Phi_t(z^t) = 1.$$

Trading in securities markets A non-zero measure of these households can trade bonds and stocks in securities markets that open every period. These households are in partition 1. We assume that the returns of these securities span Z. In other words, the payoff space is $R^{Z\times t}$ in each period t. Households in partition 2 can only trade one-period riskless discount bonds (a cash account). We use A^j to denote the menu of traded assets for households in segment $j \in \{1, 2\}$. However, none of these households can insure directly against idiosyncratic shocks ℓ_t to their labor income by selling a claim to their labor income or by trading contingent claims on these idiosyncratic shocks.

Law of One Price We assume free portfolio formation, at least for some households, and the law of one price. There exists a unique pricing kernel Π_t in the payoff space. Since there is a non-zero measure of households that trade assets that span z_t, it only depends on the aggregate shocks z_t. Formally, this pricing kernel is the projection of any candidate pricing kernel on the space of traded payoffs $\underline{X}_t = \mathcal{R}^{Z\times t}$:

$$\frac{\Pi_t}{\Pi_{t-1}} = proj(M_t | \mathcal{R}^{Z\times t}).$$

We let P_t be the arbitrage-free price of an asset with payoffs $\{D_t^i\}$:

$$P_t^i = E_t \sum_{\tau=t}^{\infty} \frac{\Pi_\tau}{\Pi_t} D_\tau^i. \tag{21}$$

for any non-negative stochastic dividend process D_t^i that is measurable w.r.t z^t.

Household Problem We adopted the approach of Cuoco and He (2001): We let agents trade a full set of Arrow securities (contingent on both aggregate and idiosyncratic shock histories), but impose *measurability* restrictions on the positions in these securities.

After collecting their labor income and their payoffs from the Arrow securities, households buy consumption in spot markets and take Arrow positions $a_{t+1}(\ell^{t+1}, z^{t+1})$ in the securities markets subject to a standard budget constraint:

$$c_t + E_t\left[\frac{\Pi_{t+1}}{\Pi_t} a_t(\ell^{t+1}, z^{t+1})\right] + \sum_{j \in A^j} P_t^j s_{t+1}^j \leq \theta_t,$$

where s denotes the shares in a security j that is in the trading set of that agent. In the second term on the left-hand side, the expectations operator arises because we sum across all states of nature tomorrow and weight the price of each of the corresponding Arrow securities by the probability of that state arising. Wealth evolves according to:

$$\theta_{t+1} = a_t(\ell^{t+1}, z^{t+1}) + \eta_{t+1} + \sum_{j \in A^j} \left[P_{t+1}^j + D_{t+1}^j\right] s_t^j,$$

subject to a measurability constraint:

$$a_t(\ell^{t+1}, z^{t+1}) \text{ is measurable w.r.t. } A_t^j(\ell^{t+1}, z^{t+1}),\ j \in \{1, 2\},$$

and subject to a generic borrowing or solvency constraint:

$$a_t(\ell^{t+1}, z^{t+1}) \geq B_t(\ell^t, z^t).$$

These measurability constraints limit the dependence of total household financial wealth on (z^{t+1}, ℓ^{t+1}). For example, for those households in partition 2 that only trade a risk-free bond, $A_t^2(\ell^{t+1}, z^{t+1}) = (\ell^t, z^t)$, because their net wealth can only depend on the history of aggregate and idiosyncratic states up until t. The households in partition 1, who do trade in stock and bond markets, can have net wealth that additionally depends on the aggregate state at time $t+1$: $A_t^1(\ell^{t+1}, z^{t+1}) = (\ell^t, z^{t+1})$.

Pricing of Household Human wealth In the absence of arbitrage opportunities, we can eliminate trade in actual securities, and the budget constraint reduces to:

$$c_t + E_t\left[\frac{\Pi_{t+1}}{\Pi_t} a_t(\ell^{t+1}, z^{t+1})\right] \leq a_{t-1}(\ell^t, z^t) + \eta_t.$$

By forward substitution of $a_t(\ell^{t+1}, z^{t+1})$ in the budget constraint, and by imposing the transversality condition on household net wealth:

$$\lim_{t \to \infty} \Pi_t a_t(\ell^t, z^t) = 0,$$

it becomes apparent that the expression for financial wealth is :

$$\begin{aligned} a_{t-1}(\ell^t, z^t) &= E_t\left[\sum_{\tau=t}^{\infty} \frac{\Pi_\tau}{\Pi_t} \left(c_\tau(\ell^\tau, z^\tau) - \eta_\tau(\ell^\tau, z^\tau)\right)\right] \\ &= E_t\left[\sum_{\tau=t}^{\infty} \frac{\Pi_\tau}{\Pi_t} c_\tau(\ell^\tau, z^\tau)\right] - E_t\left[\sum_{\tau=t}^{\infty} \frac{\Pi_\tau}{\Pi_t} \eta_\tau(\ell^\tau, z^\tau)\right] \end{aligned}$$

The equation states that non-human wealth (on the left) equals the present discounted value of consumption (total wealth) minus the present discounted value of labor income (human wealth). The value of a claim to $c-y$ is *uniquely* pinned down, because the object on the left hand side is traded financial wealth.

Pricing of Aggregate Human Wealth Let Φ_0 denote the measure at time 0 over the history of idiosyncratic shocks. The (shadow) price of a claim to aggregate labor income at time 0 is given by the aggregation of the valuation of the household labor income streams:

$$\begin{aligned} & \int E_0 \left[\sum_{t=0}^{\infty} \frac{\Pi_t}{\Pi_0} \left(\widehat{c}_t(\ell^t, z_t) C_t(z^t) - \widehat{\eta}_t(\ell^t, z_t) L_t(z^t) \right) \right] d\Phi_0 \\ = \quad & E_0 \left[\sum_{t=0}^{\infty} \frac{\Pi_t}{\Pi_0} \int \left(\widehat{c}_t(\ell^t, z_t) d\Phi_t(z^t) C_t(z^t) - \widehat{\eta}_t(\ell^t, z_t) d\Phi_t(z^t) L_t(z^t) \right) \right], \\ = \quad & E_0 \left[\sum_{t=0}^{\infty} \frac{\Pi_t}{\Pi_0} [C_t(z^t) - L_t(z^t)] \right], \end{aligned}$$

where we have used the fact that the pricing kernel Π_t does not depend on the idiosyncratic shocks, the labor income shares integrate to one $\int \widehat{\eta}_t(\ell^t, z_t) d\Phi_t(z^t) = 1$, and the consumption shares integrate to one $\int \widehat{c}_t(\ell^t, z_t) d\Phi_t(z^t) = 1$, in which $\Phi_t(z^t)$ is the distribution of household histories ℓ^t conditional on being in aggregate node z^t.

Under the maintained assumption that the traded assets span aggregate uncertainty, this implies that aggregate human wealth is the present discounted value of aggregate labor income and that total wealth is the present discounted value of aggregate consumption, and that the discounting is done with the projection of the SDF on the space of traded payoff space. Put differently, the discount factor is the same one that prices tradeable securities, such as stocks and bonds. This result follows directly from aggregating households' budget constraints. The result obtains despite the fact that human wealth is non-tradeable in this model, and therefore, markets are incomplete.

Since the above argument only relied on iterating forward on the budget constraint, we did not need to know the exact form of the equilibrium SDF. Chien, Cole, and Lustig (2007) show that the SDF in this environment depends on the evolution of the wealth distribution over time. More precisely, for each agent, one needs to keep track of a cumulative Lagrange multiplier which changes whenever the measurability constraints or the borrowing constraints bind. One cross-sectional moment of these cumulative multipliers suffices to keep track of ho the wealth distribution affects asset prices. Because of the law of large numbers, that moments only depends on the aggregate history z^t. Similar aggregation results are derived in Constantinides and Duffie (1996) and in the limited commitment models of Lustig and Chien (2007) and Lustig and Van Nieuwerburgh (2007a). To sum up, in the presence of heterogeneous agents who cannot trade idiosyncratic labor income risk, there is an additional source of aggregate risk which captures the evolution of the wealth distribution. But asset prices will fully reflect that source of aggregate risk so that our procedure remains valid in such a world.

No Spanning If the traded payoffs do not span the aggregate shocks then the preceding argument still goes through for the projection of the candidate SDF on the space of traded payoffs:

$$\frac{\Pi_t^*}{\Pi_{t-1}^*} = proj\left(M_t | \underline{X}_t\right).$$

We can still price the aggregate consumption and labor income claims using Π^*. In this case, the part of non-traded payoffs that is orthogonal to the traded payoffs, may be priced:

$$\begin{aligned} E_t \left[\left(C_{t+1} - proj\left(C_{t+1} | \underline{X}_{t+1} \right) \right) \Pi_{t+1}^* \right] & \neq 0, \\ E_t \left[\left(Y_{t+1} - proj\left(Y_{t+1} | \underline{X}_{t+1} \right) \right) \Pi_{t+1}^* \right] & \neq 0, \end{aligned}$$

where we assume that $\underline{X}$ includes a constant so that the residuals are mean zero. In the main text, we compute good-deal bounds.

B Data Appendix

B.1 Macroeconomic Series

Labor income Our data are quarterly and span the period 1952.I-2006.IV. They are compiled from the most recent data available. Labor income is computed from NIPA Table 2.1 as wage and salary disbursements (line 3) + employer contributions for employee pension and insurance funds (line 7) + government social benefits to persons (line 17) - contributions for government social insurance (line 24) + employer contributions for government social insurance (line 8) - labor taxes. As in Lettau and Ludvigson (2001a), labor taxes are defined by imputing a share of personal current taxes (line 25) to labor income, with the share calculated as the ratio of wage and salary disbursements to the sum of wage and salary disbursements, proprietors' income (line 9), and rental income of persons with capital consumption adjustment (line 12), personal interest income (line 14) and personal dividend income (line 15). The series is seasonally-adjusted at annual rates (SAAR), and we divide it by 4. Because net worth of non-corporate business and owners' equity in farm business is part of financial wealth, it cannot also be part of human wealth. Consequently, labor income excludes proprietors' income.

Consumption Non-housing consumption consists of non-housing, non-durable consumption and non-housing durable consumption. Consumption data are taken from Table 2.3.5. from the Bureau of Economic Analysis' National Income and Product Accounts (BEA, NIPA). Non-housing, non-durable consumption is measured as the sum of non-durable goods (line 6) + services (line 13) - housing services (line14).

Non-housing durable consumption is unobserved and must be constructed. From the BEA, we observe durable *expenditures*. The value of the durables (Flow of Funds, see below) at the end of two consecutive quarters and the durable expenditures allows us to measure the implicit depreciation rate that entered in the Flow of Fund's calculation. We average that depreciation rate over the sample; it is δ=5.293% per quarter. We apply that depreciation rate to the value of the durable stock at the beginning of the current period (= measured as the end of the previous quarter) to get a time-series of this period's durable *consumption*.

We use housing services consumption (BEA, NIPA, Table 2.3.5, line 14) as the dividend stream from housing wealth. The BEA measures rent for renters and imputes a rent for owners. These series are SAAR, so we divide them by 4 to get quarterly values.

Total consumption is the sum of non-housing non-durable, non-housing durable, and housing consumption.

Population and deflation Throughout, we use the disposable personal income deflator from the BEA (Table 2.1, implied by lines 36 and 37) as well as the BEA's population series (line 38).

B.2 Financial Series

Stock market return We use value-weighted quarterly returns (NYSE, AMEX, and NASDAQ) from CRSP as our measure of the stock market return. In constructing the dividend-price ratio, we use the repurchase-yield adjustment advocated by Boudoukh, Michaely, Richardson, and Roberts (2004). We also add the dividends over the current and past three quarters, so as to obtain a price-dividend ratio that is comparable with an annual number.

Additional cross-sectional stock returns We use the 25 size and value equity portfolio returns from Kenneth French. We form log real quarterly returns.

Bond yields We use the nominal yield on a 3-month Treasury bill from Fama (CRSP file) as our measure of the risk-free rate. We also use the yield spread between a 5-year Treasury note and a 3-month Treasury bill as a return predictor. The 5-year yield is obtained from the Fama-Bliss data (CRSP file). The same Fama-Bliss yields of maturities 1-, 2-, 3-, 4-, and 5-years are used to form annual forward rates and to form 1-year excess returns in the Cochrane-Piazzesi excess bond return regression.

As additional yields to fit, we use the Fama-Bliss yields at the 1-year, and 3-year maturities. We also use yields on nominal bonds at maturities 10 and 20 years. We use yield data from the Federal Reserve Bank of Saint Louis (FRED II) for the latter, and construct the spread with the 5-year yield from FRED. The 10- and 20-year yields we use are the sum of the 5-year Fama-Bliss yield and the 10-5 and 20-5 yield spread from FRED. This is to adjust for any level differences in the 5-year yield between the two data sources. The 20-year yield data are missing from 1987.I until 1993.III. The estimation can handle these missing observations because it minimizes the sum of squared differences between model-implied and observed yields, where the sum is only taken over available dates.

In order to plot the average yield curve in Figure 1, and only for this purpose, we also use the 7-5 year and the 30-5 year spread from FRED II. We add them to the 5-year yield from Fama-Bliss to form the 7-year and 30-year yield series. Since the 7-year yield data are missing from 1953.4-1969.6, we use spline interpolation (using the 1-, 2-, 5-, 10-, and 20-year yields) to fill in the missing data. The 30-year bond yield data are missing from 1953.4-1977.1 and from 2002.3-2006.1. We use the 20-year yield in those periods as a proxy. In the period where the 20-year yield is absent, we use the 30-year yield data in that period as a proxy. The resulting average 5-year yield is 6.11% per annum (straight from Fama-Bliss), the average 7-year yield is 6.25%, 10-year yield is 6.32%, 20-year is 6.51%, and the average 30-year yield is 6.45%.

C No-Arbitrage Model Details

C.1 Proof of Proposition 2

Proof. To find A_0^c and A_1^c, we need to solve the Euler equation for a claim to aggregate consumption. This Euler equation can either be thought of as the Euler equation that uses the nominal log SDF $m_{t+1}^{\$}$ to price the nominal total wealth return $\pi_{t+1} + r_{t+1}^c$ or the real log SDF $m_{t+1}^{\$} + \pi_{t+1}$ to price the real return r_{t+1}^c:

$$\begin{aligned}
1 &= E_t[\exp\{m_{t+1}^{\$} + \pi_{t+1} + r_{t+1}^c\}] \\
&= E_t[\exp\{-y_t^{\$}(1) - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + \pi_0 + e_\pi' z_{t+1} + \mu_c + e_{\Delta c}' z_{t+1} + A_0^c + A_1^{c\prime} z_{t+1} + \kappa_0^c - \kappa_1^c\left(A_0^c + A_1^{c\prime} z_t\right)\}] \\
&= \exp\{-y_0^{\$}(1) + \pi_0 - e_{yn}' z_t - \frac{1}{2}\Lambda_t'\Lambda_t + e_\pi'\Psi z_t + \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c - \kappa_1^c A_1^{c\prime} z_t + \left(e_{\Delta c}' + A_1^{c\prime}\right)\Psi z_t\} \times \\
&\quad E_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1} + (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\varepsilon_{t+1}\}\right]
\end{aligned}$$

First, note that because of log-normality of ε_{t+1}, the last line equals:

$$\exp\left\{\frac{1}{2}\left(\Lambda_t'\Lambda_t + (e_{\Delta c} + e_\pi + A_1^c)'\Sigma(e_{\Delta c} + e_\pi + A_1^c)' - 2(e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_t\right)\right\}$$

Substituting in for the expectation, as well as for the affine expression for Λ_t, we get

$$\begin{aligned} 1 &= \exp\{-y_0^{\$}(1) + \pi_0 - e_{yn}' z_t + \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c - \kappa_1^c A_1^{c\prime} z_t + (e_{\Delta c} + e_\pi + A_1^c)'\Psi z_t\} \times \\ &\quad \exp\{\frac{1}{2}(e_{\Delta c} + e_\pi + A_1^c)'\Sigma(e_{\Delta c} + e_\pi + A_1^c) - (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}(\Lambda_0 + \Lambda_1 z_t)\} \end{aligned}$$

Taking logs on both sides, an collecting the constant terms and the terms in z, we obtain the following:

$$\begin{aligned} 0 &= \{-y_0^{\$}(1) + \pi_0 + \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c + \frac{1}{2}(e_{\Delta c} + e_\pi + A_1^c)'\Sigma(e_{\Delta c} + e_\pi + A_1^c) - (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_0\} + \\ &\quad \{-e_{yn}' - \kappa_1^c A_1^{c\prime} + (e_{\Delta c} + e_\pi + A_1^c)'\Psi - (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_1\} z_t \end{aligned}$$

This equality needs to hold for all z_t. This is a system of $N+1$ equations in $N+1$ unknowns:

$$0 = -y_0^{\$}(1) + \pi_0 + \kappa_0^c + (1-\kappa_1^c)A_0^c + \mu_c + \frac{1}{2}(e_{\Delta c} + e_\pi + A_1^c)'\Sigma(e_{\Delta c} + e_\pi + A_1^c) - (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_0, \tag{22}$$

$$0 = (e_{\Delta c} + e_\pi + A_1^c)'\Psi - \kappa_1^c A_1^{c\prime} - e_{yn}' - (e_{\Delta c} + e_\pi + A_1^c)'\Sigma^{\frac{1}{2}}\Lambda_1. \tag{23}$$

The real short yield $y_t(1)$, or *risk-free rate*, satisfies $E_t[\exp\{m_{t+1} + y_t(1)\}] = 1$. Solving out this Euler equation, we get:

$$\begin{aligned} y_t(1) &= y_t^{\$}(1) - E_t[\pi_{t+1}] - \frac{1}{2}e_\pi'\Sigma e_\pi + e_\pi'\Sigma^{\frac{1}{2}}\Lambda_t \\ &= y_0(1) + \left[e_{yn}' - e_\pi'\Psi + e_\pi'\Sigma^{\frac{1}{2}}\Lambda_1\right] z_t && (24) \\ y_0(1) &\equiv y_0^{\$}(1) - \pi_0 - \frac{1}{2}e_\pi'\Sigma e_\pi + e_\pi'\Sigma^{\frac{1}{2}}\Lambda_0 && (25) \end{aligned}$$

The real short yield is the nominal short yield minus expected inflation minus a Jensen adjustment minus the inflation risk premium. Using the expression (25) for $y_0(1)$ in equation (22) delivers equation (6) in the main text. Likewise, using (24) in equation (23) delivers equation (7). □

Corollary 5. *The log price-dividend ratio on human wealth is a linear function of the (demeaned) state vector z_t*

$$pd_t^l = A_0^l + A_1^l z_t$$

where the following recursions pin down A_0^l and A_1^l:

$$\begin{aligned} 0 &= \kappa_0^l + (1-\kappa_1^l)A_0^l + \mu_l - y_0(1) + \frac{1}{2}(e_{\Delta l}' + A_1^{l\prime})\Sigma(e_{\Delta l} + A_1^l) - (e_{\Delta l}' + A_1^{l\prime})\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right), \\ 0 &= \left(e_{\Delta l} + e_\pi + A_1^l\right)'\Psi - \kappa_1^l A_1^{l\prime} - e_{yn}' - (e_{\Delta l} + e_\pi + A_1^l)'\Sigma^{\frac{1}{2}}\Lambda_1. \end{aligned}$$

The proof is identical to the proof of Proposition 2, and obtains by replacing μ_c by μ_l and the selector vector e_c by $e_{\Delta l}$. The linearization constants κ_0^l and κ_1^l relate to A_0^l through the analog of equation (3).

C.2 Campbell-Shiller Variance Decomposition

By iterating forward on the total wealth return equation (2), we can link the log wealth-consumption ratio at time t to expected future total wealth returns and consumption growth rates:

$$wc_t = \frac{\kappa_0^c}{\kappa_1^c - 1} + \sum_{j=1}^{H} (\kappa_1^c)^{-j} \Delta c_{t+j} - \sum_{j=1}^{H} (\kappa_1^c)^{-j} r_{t+j}^c + (\kappa_1^c)^{-H} wc_{t+H}. \tag{26}$$

Because this expression holds both ex-ante and ex-post, one is allowed to add the expectation sign on the right-hand side. Imposing the transversality condition as $H \to \infty$ kills the last term, and delivers the familiar Campbell and Shiller (1988) decomposition for the price-dividend ratio of the consumption claim:

$$wc_t = \frac{\kappa_0^c}{\kappa_1^c - 1} + E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j} \Delta c_{t+j}\right] - E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j} r_{t+j}\right] = \frac{\kappa_0^c}{\kappa_1^c - 1} + \Delta c_t^H - r_t^H. \tag{27}$$

where the second equality follows from the definitions

$$\Delta c_t^H \equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j} \Delta c_{t+j}\right] = e_c' \Psi (\kappa_1^c I - \Psi)^{-1} z_t, \tag{28}$$

$$r_t^H \equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j} r_{t+j}\right] = \left[(e_c + A_1^c)'\Psi - \kappa_1^c A_1^{c\prime}\right] (\kappa_1^c I - \Psi)^{-1} z_t, \tag{29}$$

where I is the $N \times N$ identity matrix. The first equation for the cash-flow component Δc_t^H follows from the VAR dynamics, while the second equation for the discount rate component r_t^H follows from Proposition 2 and the definition of the total wealth return equation (2).

Using expressions (29) and (28) and the log-linearity of the wealth-consumption ratio, we obtain analytical expressions for the following variance and covariance terms:

$$V[wc_t] = A_1^{c\prime} \Omega A_1^c \tag{30}$$

$$Cov\left[wc_t, \Delta c_t^H\right] = A_1^{c\prime} \Omega (\kappa_1^c I - \Psi')^{-1} \Psi' e_c \tag{31}$$

$$Cov\left[wc_t, -r_t^H\right] = A_1^{c\prime} \Omega \left[A_1^c - (\kappa_1^c I - \Psi)^{-1} \Psi' e_c'\right] \tag{32}$$

$$V\left[\Delta c_t^H\right] = e_c' \Psi (\kappa_1^c I - \Psi)^{-1} \Omega (\kappa_1^c I - \Psi')^{-1} \Psi' e_c \tag{33}$$

$$V\left[r_t^H\right] = \left[(e_c' + A_1^{c\prime})\Psi - \kappa_1^c A_1^{c\prime}\right] (\kappa_1^c I - \Psi)^{-1} \Omega (\kappa_1^c I - \Psi')^{-1} \left[\Psi'(e_c + A_1^c) - \kappa_1^c A_1^c\right] \tag{34}$$

$$Cov\left[r_t^H, \Delta c_t^H\right] = \left[(e_c' + A_1^{c\prime})\Psi - \kappa_1^c A_1^{c\prime}\right] (\kappa_1^c I - \Psi)^{-1} \Omega (\kappa_1^c I - \Psi')^{-1} \Psi' e_c \tag{35}$$

where $\Omega = E[z_t' z_t]$ is the second moment matrix of the state z_t.

C.3 Nominal Term Structure

Proposition 6. *Nominal bond yields are affine in the state vector:*

$$y_t^\$(\tau) = -\frac{A^\$(\tau)}{\tau} - \frac{B^\$(\tau)'}{\tau} z_t,$$

where the coefficients $A^\$(\tau)$ and $B^\$(\tau)$ satisfy the following recursions

$$
\begin{aligned}
A^\$(\tau+1) &= -y_0^\$(1) + A^\$(\tau) + \frac{1}{2}\left(B^\$(\tau)\right)'\Sigma\left(B^\$(\tau)\right) - \left(B^\$(\tau)\right)'\Sigma^{\frac{1}{2}}\Lambda_0, && (36)\\
\left(B^\$(\tau+1)\right)' &= \left(B^\$(\tau)\right)'\Psi - e_{yn}' - \left(B^\$(\tau)\right)'\Sigma^{\frac{1}{2}}\Lambda_1, && (37)
\end{aligned}
$$

initialized at $A^\$(0) = 0$ and $B^\$(0) = \mathbf{0}$.

Proof. We conjecture that the $t+1$-price of a τ-period bond is exponentially affine in the state

$$\log(P_{t+1}^\$(\tau)) = A^\$(\tau) + \left(B^\$(\tau)\right)' z_{t+1}$$

and solve for the coefficients $A^\$(\tau+1)$ and $B^\$(\tau+1)$ in the process of verifying this conjecture using the Euler equation:

$$
\begin{aligned}
P_t^\$(\tau+1) &= E_t[\exp\{m_{t+1}^\$ + \log\left(P_{t+1}^\$(\tau)\right)\}]\\
&= E_t[\exp\{-y_t^\$(1) - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + A^\$(\tau) + \left(B^\$(\tau)\right)' z_{t+1}\}]\\
&= \exp\{-y_0^\$(1) - e_{yn}'z_t - \frac{1}{2}\Lambda_t'\Lambda_t + A^\$(\tau) + \left(B^\$(\tau)\right)'\Psi z_t\} \times\\
&\quad E_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1} + \left(B^\$(\tau)\right)'\Sigma^{\frac{1}{2}}\varepsilon_{t+1}\}\right]
\end{aligned}
$$

We use the log-normality of ε_{t+1} and substitute for the affine expression for Λ_t to get:

$$P_t^\$(\tau+1) = \exp\{-y_0^\$(1) - e_{yn}'z_t + A^\$(\tau) + \left(B^\$(\tau)\right)'\Psi z_t + \frac{1}{2}\left(B^\$(\tau)\right)'\Sigma\left(B^\$(\tau)\right) - \left(B^\$(\tau)\right)'\Sigma^{\frac{1}{2}}(\Lambda_0 + \Lambda_1 z_t)\}$$

Taking logs and collecting terms, we obtain a linear equation for $\log(p_t(\tau+1))$:

$$\log\left(P_t^\$(\tau+1)\right) = A^\$(\tau+1) + \left(B^\$(\tau+1)\right)' z_t,$$

where $A^\$(\tau+1)$ satisfies (36) and $B^\$(\tau+1)$ satisfies (37). □

Real bond yields, $y_t(\tau)$, denoted without the \$ superscript, are affine as well with coefficients that follow similar recursions:

$$
\begin{aligned}
A(\tau+1) &= -y_0(1) + A(\tau) + \frac{1}{2}(B(\tau))'\Sigma(B(\tau)) - (B(\tau))'\Sigma^{\frac{1}{2}}\left(\Lambda_0 - \Sigma^{\frac{1}{2}\prime}e_\pi\right),\\
(B(\tau+1))' &= (e_\pi + B(\tau))'\Psi - e_{yn}' - (e_\pi + B(\tau))'\Sigma^{\frac{1}{2}}\Lambda_1.
\end{aligned}
$$

For $\tau = 1$, we recover the expression for the risk-free rate in (24)-(25).

C.4 Dividend Strips

We define the return on equity conform the literature as $R_{t+1}^m = \frac{P_{t+1}^m + D_{t+1}^m}{P_t^m}$, where P_t^m is the end-of-period price on the equity market. A log-linearization delivers:

$$r_{t+1}^m = \kappa_0^m + \Delta d_{t+1}^m + \kappa_1^m pd_{t+1}^m - pd_t^m. \quad (38)$$

The unconditional mean stock return is $r_0^m = \kappa_0^m + (\kappa_1^m - 1)A_0^m + \mu_m$, where $A_0^m = E[pd_t^m]$ is the unconditional average log price-dividend ratio on equity and $\mu_m = E[\Delta d_t^m]$ is the unconditional mean dividend growth rate. The linearization constants κ_0^m and κ_1^m are different from the other wealth concepts because the timing of the return is different:

$$\kappa_1^m = \frac{e^{A_0^m}}{e^{A_0^m}+1} < 1 \text{ and } \kappa_0^m = \log\left(e^{A_0^m}+1\right) - \frac{e^{A_0^m}}{e^{A_0^m}+1}A_0^m. \tag{39}$$

Even though these constants arise from a linearization, we *define* log dividend growth so that the return equation holds exactly, given the CRSP series for $\{r_t^m, pd_t^m\}$. Our state vector z contains the (demeaned) return on the stock market, $r_{t+1}^m - r_0^m$, and the (demeaned) log price-dividend ratio $pd^m - A_0^m$. The definition of log equity returns allows us to back out dividend growth:

$$\Delta d_{t+1}^m = \mu^m + \left[(e_{rm} - \kappa_1^m e_{pd})'\Psi + e_{pd}'\right] z_t + (e_{rm} - \kappa_1^m e_{pd})' \Sigma^{\frac{1}{2}} \varepsilon_{t+1}$$

Proposition 7. *Log price-dividend ratios on dividend strips are affine in the state vector:*

$$p_t^d(\tau) = A^m(\tau) + B^{m\prime}(\tau) z_t,$$

where the coefficients $A^m(\tau)$ and $B^m(\tau)$ follow recursions

$$\begin{aligned}
A^m(\tau+1) &= A^m(\tau) + \mu_m - y_0(1) + \frac{1}{2}\left(e_{rm} - \kappa_1^m e_{pdm} + B^c(\tau)\right)' \Sigma \left(e_{rm} - \kappa_1^m e_{pdm} + B^m(\tau)\right) \\
&\quad - \left(e_{rm} - \kappa_1^m e_{pdm} + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}} e_\pi\right), \\
B^m(\tau+1)' &= \left(e_{rm} - \kappa_1^m e_{pdm} + e_\pi + B^m(\tau)\right)' \Psi + e_{pdm}' - e_{yn}' - \left(e_{rm} - \kappa_1^m e_{pdm} + e_\pi + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} \Lambda_1,
\end{aligned}$$

initialized at $A^m(0) = 0$ and $B^m(0) = 0$.

Proof. We conjecture that the log $t+1$-price of a τ-period strip, scaled by the dividend in period $t+1$, is affine in the state

$$p_{t+1}^d(\tau) = \log\left(P_{t+1}^d(\tau)\right) = A^m(\tau) + B^m(\tau)' z_{t+1}$$

and solve for the coefficients $A^m(\tau+1)$ and $B^m(\tau+1)$ in the process of verifying this conjecture using the Euler equation:

$$\begin{aligned}
P_t^d(\tau+1) &= E_t[\exp\{m_{t+1}^\$ + \pi_{t+1} + \Delta d_{t+1}^m + \log\left(p_{t+1}^m(\tau)\right)\}] \\
&= E_t[\exp\{-y_t^\$(1) - \frac{1}{2}\Lambda_t'\Lambda_t - \Lambda_t'\varepsilon_{t+1} + \pi_0 + e_\pi' z_{t+1} + \Delta d_{t+1}^m + A^m(\tau) + B^m(\tau)' z_{t+1}\}] \\
&= \exp\{-y_0^\$(1) - e_{yn}' z_t - \frac{1}{2}\Lambda_t'\Lambda_t + \pi_0 + e_\pi'\Psi z_t + \mu_m + \left[(e_{rm} - \kappa_1^m e_{pd})'\Psi + e_{pd}'\right] z_t + A^m(\tau) + B^m(\tau)'\Psi z_t\} \times \\
&\quad E_t\left[\exp\{-\Lambda_t'\varepsilon_{t+1} + \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} \varepsilon_{t+1}\right]
\end{aligned}$$

We use the log-normality of ε_{t+1} and substitute for the affine expression for Λ_t to get:

$$\begin{aligned}
P_t^d(\tau+1) &= \exp\{-y_0^\$(1) - e_{yn}' z_t + \pi_0 + \mu_m + A^m(\tau) + \left[\left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Psi + e_{pd}'\right] z_t + \\
&\quad \frac{1}{2}\left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Sigma \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right) - \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} (\Lambda_0 + \Lambda_1 z_t)\}
\end{aligned}$$

Taking logs and collecting terms, we obtain a log-linear expression for $p_t^m(\tau+1)$:

$$p_t^d(\tau+1) = A^m(\tau+1) + B^m(\tau+1)' z_t,$$

where

$$
\begin{aligned}
A^m(\tau+1) &= A^m(\tau) + \mu_m - y_0^{\$}(1) + \pi_0 + \frac{1}{2}\left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^c(\tau)\right)' \Sigma \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right) \\
&\quad - \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} \Lambda_0, \\
B^m(\tau+1)' &= \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Psi + e_{pd}' - e_{yn}' - \left(e_{rm} - \kappa_1^m e_{pd} + e_\pi + B^m(\tau)\right)' \Sigma^{\frac{1}{2}} \Lambda_1
\end{aligned}
$$

We recover the recursions of Proposition 7 after substituting out the expressions for the nominal yields using equations (24) and (25). □

C.5 Consumption Strips

Proposition 8. *Log price-dividend ratios on consumption strips are affine in the state vector:*

$$p_t^c(\tau) = A^c(\tau) + B^{c\prime}(\tau) z_t,$$

where the coefficients $A^c(\tau)$ and $B^c(\tau)$ follow recursions:

$$
\begin{aligned}
A^c(\tau+1) &= A^c(\tau) + \mu_c - y_0(1) + \frac{1}{2}\left(e_c + B^c(\tau)\right)' \Sigma \left(e_c + B^c(\tau)\right) - \left(e_c + B^c(\tau)\right)' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}} e_\pi\right), \\
B^c(\tau+1)' &= \left(e_c + e_\pi + B^c(\tau)\right)' \Psi - e_{yn}' - \left(e_c + e_\pi + B^c(\tau)\right)' \Sigma^{\frac{1}{2}} \Lambda_1.
\end{aligned}
$$

initialized at $A^c(0) = 0$ and $B^c(0) = 0$.

The proof is identical to that of Proposition 7.

We decompose the yield on a τ-period consumption strip from Proposition 8, $y_t^c(\tau) = -p_t^c(\tau)/\tau$, into the yield on a τ-period real bond adjusted for consumption growth plus the yield on the consumption cash-flow risk security $y_t^{ccr}(\tau)$:

$$y_t^c(\tau) = (y_t(\tau) - \mu_c) + y_t^{ccr}(\tau).$$

The former can be thought of as the period-τ coupon yield on a real perpetuity with cash-flows that grow at a deterministic rate μ_c, while the latter captures the cash-flow risk in the consumption claim. We have that $y_t^{ccr}(\tau) = -p_t^{ccr}(\tau)/\tau$. Since the log price-dividend ratio of the consumption strips and the log real bond prices are both affine, so is the log price-dividend ratio of the consumption cash-flow risk security: $\log p_t^{ccr}(\tau) = A^{ccr}(\tau) + B^{ccr}(\tau) z_t$. It is easy to show that its coefficients follow the recursions:

$$
\begin{aligned}
A^{ccr}(\tau+1) &= A^{ccr}(\tau) + \frac{1}{2}\left(e_c + B^{ccr}(\tau)\right)' \Sigma \left(e_c + B^{ccr}(\tau)\right) + \left(e_c + B^{ccr}(\tau)\right)' \Sigma B(\tau) \\
&- \left(e_c + B^{ccr}(\tau)\right)' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}} e_\pi\right), \\
B^{ccr}(\tau+1)' &= \left(e_c + B^{ccr}(\tau)\right)' \Psi - \left(e_c + B^{ccr}(\tau)\right)' \Sigma^{\frac{1}{2}} \Lambda_1.
\end{aligned}
$$

C.6 Point Estimates

Below, we report the point estimates for the VAR companion matrix Ψ, the Cholesky decomposition of the covariance matrix $\Sigma^{.5}$ (multiplied by 100), and the market price of risk parameters Λ_0 and Λ_1 for our benchmark specification. We recall that the market price of risk parameter matrix Λ_1 pre-multiplies the state z_t, which has a (non-standardized) covariance matrix Ω.

$$
\Psi \quad = \quad \begin{bmatrix}
\mathbf{0.3053} & \mathbf{1.2456} & \mathbf{-0.4565} & \mathbf{3.3037} & 0 & 0 & 0 & 0 & 0 & 0 \\
-0.0035 & \mathbf{0.8965} & \mathbf{0.1022} & 0.1590 & 0 & 0 & 0 & 0 & 0 & 0 \\
\mathbf{-0.0371} & \mathbf{0.1690} & \mathbf{0.7312} & 0.0930 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.0086 & 0.0289 & -0.0424 & \mathbf{0.6973} & 0 & 0 & 0 & 0 & 0 & 0 \\
0.3201 & -2.3434 & -0.9365 & -2.8085 & \mathbf{0.9131} & -0.0230 & 0 & 0 & 0 & 0 \\
0.2257 & -1.8392 & -1.5961 & -0.9984 & \mathbf{-0.0890} & 0.0429 & 0 & 0 & 0 & 0 \\
0.0139 & 0.0466 & -0.5851 & 0.6283 & -0.0105 & \mathbf{0.0708} & 0 & 0 & 0 & 0 \\
0.1506 & -0.2318 & -0.6371 & -0.0617 & -0.0130 & 0.0769 & 0 & 0 & 0 & 0 \\
0.0251 & -0.0202 & -0.0557 & 0.0350 & 0.0006 & -0.0020 & -0.0002 & 0.0068 & \mathbf{0.3014} & 0 \\
\mathbf{0.1032} & -0.2691 & 0.0091 & -0.6800 & 0.0016 & -0.0064 & 0.0495 & -0.0207 & \mathbf{0.5887} & -0.1390
\end{bmatrix}
$$

$$
\Sigma^{.5} \times 100 \quad = \quad \begin{bmatrix}
1.4427 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
-0.0714 & 0.2331 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
-0.0519 & 0.0352 & 0.2998 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.0838 & -0.1071 & 0.0136 & 0.1012 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.7999 & -0.6542 & -1.6896 & -0.3572 & 7.8117 & 0 & 0 & 0 & 0 & 0 \\
0.6247 & -0.6375 & -1.8040 & -0.4144 & 7.3491 & 2.5131 & 0 & 0 & 0 & 0 \\
0.4139 & -0.1845 & -0.4411 & -0.1536 & 1.7789 & 0.5328 & 2.5774 & 0 & 0 & 0 \\
1.0101 & -0.4506 & -0.6907 & -0.1438 & 2.2836 & 0.8748 & 3.8193 & 4.7776 & 0 & 0 \\
0.0292 & -0.0569 & 0.0166 & 0.0106 & 0.0801 & 0.0593 & 0.1280 & -0.0122 & 0.3631 & 0 \\
0.1751 & -0.0484 & -0.0046 & -0.0113 & 0.1342 & -0.0335 & 0.2257 & 0.2664 & 0.3247 & 0.6503
\end{bmatrix}
$$

$$
\Lambda_0' \quad = \quad \begin{bmatrix} 0 & \mathbf{-0.1781} & 0 & \mathbf{0.0309} & 0 & \mathbf{0.6601} & -0.0712 & 0.1225 & 0 & 0 \end{bmatrix}
$$

$$
\Lambda_1 \quad = \quad \begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\mathbf{4.1229} & \mathbf{-70.9093} & \mathbf{55.3549} & \mathbf{-96.4021} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\mathbf{-49.3448} & \mathbf{131.6779} & \mathbf{-48.3828} & \mathbf{249.5507} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.3578 & -101.5409 & -29.1252 & -17.9823 & \mathbf{-3.5997} & 1.7053 & 0 & 0 & 0 & 0 \\
-3.7005 & -5.7382 & 12.0120 & 41.0278 & 0.4508 & \mathbf{2.3711} & 0 & 0 & 0 & 0 \\
4.2683 & -1.2674 & 1.0349 & -29.7221 & 0.1894 & -0.5990 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
$$

We compute OLS standard errors for the elements of Ψ and report the coefficients with a t-stat greater than 1.98 in bold. Bootstrap standard errors on the market price of risk parameters are available upon request. They are derived as part of the method explained in Appendix C.7. Market price of risk estimates with bootstrap t-stats above 1.98 are indicated in bold.

C.7 Bootstrap Standard Errors

We obtain standard errors on the moments of the estimated wealth-consumption ratio by bootstrap. More precisely, we conduct two bootstrap exercises leading to two sets of standard errors. In each exercise, we draw with replacement from the VAR innovations ε_t. We draw row-by-row in order to preserve the cross-correlation structure between the state innovations (Step 1). Given the point estimates for Ψ and Σ as well as the mean vector μ, we recursively reconstruct the state vector (Step 2). We then re-estimate the mean vector, companion matrix, and innovation covariance matrix (Step 3). With the new state vector and the new VAR parameters in hand, we re-estimate the market price of risk parameters in Λ_0 and Λ_1 (Step 4). Just as in the main estimation, we use 2500 quarters to approximate the infinite-horizon sums in the strip price-dividend ratio calculations. We limit the estimation in Step 4 to 500 function evaluations for computational reasons. In some of the bootstrap iterations, the optimization in Step 4 does not find a feasible solution. This happens, for example when no parameter choices keep the human wealth share less than hundred percent or the consumption or labor income claim finite. We discard these bootstrap iterations. These new market price of risk parameters deliver a new wealth-consumption ratio time series (Step 5). With the bootstrap time series for consumption growth and the wealth-consumption ratio, we can form all the moments in Table 1. We repeat this procedure 1,000 times and report the standard deviation across the bootstrap iterations. We conduct two variations on the above algorithm. Each bootstrap exercise takes about 12 hours to compute on an 8-processor computer. The more conservative standard errors from the second bootstrap exercise are the ones reported in Table 1.

In the first exercise, we only consider sampling uncertainty in the last four elements of the state: the two factor mimicking portfolios, consumption growth, and labor income growth. We assume that all the other variables are observed without error. The idea is that national account aggregates are measured much less precisely than traded stocks and bonds. This procedure takes into account sampling uncertainty in consumption growth and its correlations with yields and with the aggregate stock market. Given our goal of obtaining standard errors around the moments of the wealth-consumption ratio, this seems like a natural first exercise. The second column of Table 4 reports the standard errors from this bootstrap exercise in parentheses. For completeness, it also reports the mean across bootstrap iterations.

In a second estimation exercise, we also consider sampling uncertainty in the first six state variables (yields and stock prices). Redrawing the yields that enter in the state space (the 1-quarter yield and the 20-1 quarter yield spread) requires also redrawing the additional yields that are used in estimation (the 4-, 12-, 40-, and 80-quarter yields) and in the formation of the Cochrane-Piazzesi factor (the 4-, 8-, 12-, and 16- quarter yields). Otherwise, the bootstrapped time-series for the yields in the state space would be disconnected from the other yields. For this second exercise, we augment the VAR with the following yield spreads: 4-20, 8-20, 12-20, 16-20, 40-20, and 80-20 quarter yield spreads. We let these spreads depend on their own lag and on the lagged 1-quarter yield. Additional dependence on the lagged 20-1 quarter yield makes little difference. In Step 1, we draw from the yield spreads-augmented VAR innovations. This allows us to take into account the cross-dependencies between all the yields in the yield curve. In addition to recursively rebuilding the state variables in Step 2, we also rebuild the six yield spreads. With the bootstrapped yields, we reconstruct the forward rates, one-year excess bond returns, re-estimate the excess bond return regression, and re-construct the Cochrane-Piazzesi factor. Steps 3 through 5 are the same as in the first exercise. One additional complication arises because the bootstrapped yields often turn negative for one or more periods. Since negative nominal yields never happen in the data and make no economic sense, we discard these bootstrap iterations. We redraw from the VAR innovations until we have 1,000 bootstrap samples with strictly positive yields at all maturities. This is akin to a rejection-sampling procedure. One drawback is that there is an upward bias in the yield curve. The average one-quarter yield is 5.13% per annum in our data sample, while it is 5.93% in our bootstrap sample. This translates in a small downward bias in the average wealth consumption ratio:

the average wealth-consumption ratio is 5.86 in the data and 5.69 in the bootstrap. The third column of Table 4 reports the standard errors from this bootstrap exercise in parentheses. As expected, the standard errors from the second bootstrap exercise are somewhat bigger. However, their difference is small.

Table 4: Bootstrap Standard Errors

This table displays bootstrap standard errors on the unconditional moments of the log wealth-consumption ratio wc, its first difference Δwc, and the log total wealth return r^c. The last but one row reports the time-series average of the conditional consumption risk premium, $E[E_t[r_t^{c,e}]]$, where $r^{c,e}$ denotes the expected log return on total wealth in excess of the risk-free rate and corrected for a Jensen term. The first column repeats the point estimates from the main text (last column of Table 1). The second and third columns report the results from two bootstrap exercises, described above. The table reports the mean and standard deviation (in parentheses) across 1,000 bootstrap iterations.

Moments	Point Estimate	Bootstrap 1	Bootstrap 2
$Std[wc]$	17.24%	16.26%	15.24%
(s.e.)		(3.39)	(4.30)
$AC(1)[wc]$	.96	.95	.93
(s.e.)		(.00)	(.03)
$AC(4)[wc]$	.85	.83	.74
(s.e.)		(.01)	(.08)
$Std[\Delta wc]$	4.86%	4.86%	5.07%
(s.e.)		(0.98)	(1.16)
$Std[\Delta c]$	.44%	.44%	.44%
(s.e.)		(.03)	(.03)
$Corr[\Delta c, \Delta wc]$	.11	.02	.12
(s.e.)		(.06)	(.06)
$Std[r^c]$	4.94%	4.90%	5.16%
(s.e.)		(2.21)	(1.16)
$Corr[r^c, \Delta c]$	.19	.12	.21
(s.e.)		(.07)	(.07)
$E[E_t[r_t^{c,e}]]$	.54%	.46%	0.53%
(s.e.)		(.11)	(.16)
$E[wc]$	5.86	6.29	5.69
(s.e.)		(.48)	(.49)

D The Long-Run Risk Model

D.1 Preferences

Let $V_t(C_t)$ denote the utility derived from consuming C_t, then the value function of the representative agent takes the following recursive form:

$$V_t(C_t) = \left[(1-\beta)C_t^{1-\rho} + \beta(\mathcal{R}_t V_{t+1})^{1-\rho}\right]^{\frac{1}{1-\rho}}, \tag{40}$$

where the risk-adjusted expectation operator is defined as:

$$\mathcal{R}_t V_{t+1} = \left(E_t V_{t+1}^{1-\alpha}\right)^{\frac{1}{1-\alpha}}.$$

For these preferences, α governs risk aversion and ρ governs the willingness to substitute consumption inter-temporally. It is the inverse of the inter-temporal elasticity of substitution. In the special case where $\rho = \alpha$, they collapse to the standard power utility preferences, used in Breeden (1979) and Lucas (1978). Epstein and Zin (1989) show that the stochastic discount factor can be written as:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\rho} \left(\frac{V_{t+1}}{\mathcal{R}_t V_{t+1}}\right)^{\rho-\alpha} \tag{41}$$

The next proposition shows that the ability to write the SDF in the long-run risk model as a function of consumption growth and the wealth-consumption ratio is general. It does not depend on the linearization of returns, nor on the assumptions on the stochastic process for consumption growth (see Hansen, Heaton, Lee, and Roussanov (2008)).

Proposition 9. *The log SDF in the non-linear version of the long-run risk model can be stated as*

$$m_{t+1} = \frac{1-\alpha}{1-\rho}\log\beta - \alpha\Delta c_{t+1} + \frac{\rho-\alpha}{1-\rho}\log\left(\frac{e^{wc_{t+1}}}{e^{wc_t}-1}\right) \tag{42}$$

Proof. We start from the value function definition in equation (40) and raise both sides to the power $1-\rho$, and subsequently divide through by $(1-\beta)C_t^{-\rho}$ to obtain:

$$\frac{V_t^{1-\rho}}{(1-\beta)C_t^{-\rho}} = C_t + \beta\frac{\left(E_t V_{t+1}^{1-\alpha}\right)^{\frac{1-\rho}{1-\alpha}}}{(1-\beta)C_t^{-\rho}} \tag{43}$$

Some algebra and the definition of the risk-adjusted expectation operator imply that

$$E_t(V_{t+1}^{1-\alpha})^{\frac{1-\rho}{1-\alpha}} = E_t(V_{t+1}^{1-\alpha})^{1-\frac{\rho-\alpha}{1-\alpha}} = \frac{E_t(V_{t+1}^{1-\alpha})}{E_t(V_{t+1}^{1-\alpha})^{\frac{\rho-\alpha}{1-\alpha}}} = \frac{E_t(V_{t+1}^{1-\alpha})}{(\mathcal{R}_t V_{t+1})^{\rho-\alpha}} = E_t\left[\frac{V_{t+1}^{\rho-\alpha}V_{t+1}^{1-\rho}}{(\mathcal{R}_t V_{t+1})^{\rho-\alpha}}\right]$$

The left-hand side shows up in equation (43), and we substitute it for the right-hand side of the equation above. Multiplying and dividing inside the expectation operator by $C_{t+1}^{-\rho}$, we get:

$$\frac{V_t^{1-\rho}}{(1-\beta)C_t^{-\rho}} = C_t + E_t\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\rho}\left(\frac{V_{t+1}}{\mathcal{R}_t V_{t+1}}\right)^{\rho-\alpha}\frac{V_{t+1}^{1-\rho}}{(1-\beta)C_{t+1}^{-\rho}}\right]$$

Note that the first three terms inside the expectation are equal to the stochastic discount factor in equation (41). This is a no-arbitrage asset pricing equation of an asset with dividend equal to aggregate consumption. The price of this asset is W_t. Hence,

$$W_t = C_t + E_t\left[M_{t+1}W_{t+1}\right] \quad \text{and} \quad W_t = \frac{V_t^{1-\rho}}{(1-\beta)C_t^{-\rho}}. \tag{44}$$

This equation, together with $E[M_{t+1}R_{t+1}^c] = 1$ delivers the return on the total wealth portfolio:

$$R_{t+1}^c = \frac{W_{t+1}}{(W_t - C_t)} = \frac{\frac{V_{t+1}^{1-\rho}}{(1-\beta)C_{t+1}^{-\rho}}}{\frac{V_t^{1-\rho}}{(1-\beta)C_t^{-\rho}} - C_t} = \frac{\frac{V_{t+1}^{1-\rho}}{(1-\beta)C_{t+1}^{-\rho}}}{\frac{\beta}{(1-\beta)C_t^{-\rho}}(\mathcal{R}_t V_{t+1})^{1-\rho}} = \beta^{-1}\left(\frac{C_{t+1}}{C_t}\right)^{\rho}\left(\frac{V_{t+1}}{\mathcal{R}_t V_{t+1}}\right)^{1-\rho}, \tag{45}$$

where the first equality is the return definition, the second one follows from the definition of W_t, and the third equality from the Bellman equation (40).

One then rearranges this equation (after raising both sides to the power $\frac{\rho-\alpha}{1-\rho}$) to get a relationship between the change in continuation values and the total wealth return:

$$\left(\frac{V_{t+1}}{\mathcal{R}_t V_{t+1}}\right)^{\rho-\alpha} = \beta^{\frac{\rho-\alpha}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\rho\frac{\rho-\alpha}{1-\rho}} \left(R^c_{t+1}\right)^{\frac{\rho-\alpha}{1-\rho}}.$$

Finally, we substitute it into the stochastic discount factor expression (41) to obtain an expression that depends only on consumption growth and the return to the wealth portfolio:

$$M_{t+1} = \beta^{\frac{1-\alpha}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\rho\frac{1-\alpha}{1-\rho}} \left(R^c_{t+1}\right)^{\frac{\rho-\alpha}{1-\rho}} \tag{46}$$

Instead, we rewrite the return on the wealth portfolio in terms of the wealth-consumption ratio WC

$$R^c_{t+1} = \frac{WC_{t+1}}{WC_t - 1}\frac{C_{t+1}}{C_t} = \frac{e^{wc_{t+1}}}{e^{wc_t} - 1}\frac{C_{t+1}}{C_t},$$

Substituting this into equation (46) delivers an expression that depends only on consumption growth and the wealth-consumption ratio:

$$M_{t+1} = \beta^{\frac{1-\alpha}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\alpha} \left(\frac{WC_{t+1}}{WC_t - 1}\right)^{\frac{\rho-\alpha}{1-\rho}} = \beta^{\frac{1-\alpha}{1-\rho}} \left(\frac{C_{t+1}}{C_t}\right)^{-\alpha} \left(\frac{e^{wc_{t+1}}}{e^{wc_t} - 1}\right)^{\frac{\rho-\alpha}{1-\rho}} \tag{47}$$

□

The log SDF expression in the Bansal and Yaron (2004) model is a special case of this general, non-linear formulation. Indeed, one recovers the log SDF equation (48) below by using a first-order Taylor approximation of wc_t in equation (42) around A^c_0.

$$m_{t+1} = \frac{1-\alpha}{1-\rho}\left[\log\beta + \kappa^c_0\right] - \kappa^c_0 - \alpha\Delta c_{t+1} - \frac{\alpha-\rho}{1-\rho}\left(wc_{t+1} - \kappa^c_1 wc_t\right) \tag{48}$$

This expression follows immediately from the results in Appendix D.2 below. Since κ^c_1 turns out to be essentially 1 under the Bansal and Yaron (2004) calibration, the second asset pricing factor in the SDF is essentially the log change in the wealth-consumption ratio. A second special case obtains by approximating the last term in (42) using a first-order Taylor expansion of wc_{t+1} around wc_t instead. In that case, we obtain a three-factor model:

$$m_{t+1} \approx \frac{1-\alpha}{1-\rho}\log\beta - \alpha\Delta c_{t+1} + \frac{\rho-\alpha}{1-\rho}\log\left(\frac{e^{wc_t}}{e^{wc_t} - 1}\right) + \frac{\rho-\alpha}{1-\rho}\Delta wc_{t+1}. \tag{49}$$

Expressions (49) and (48) are functionally similar because κ^c_1 is close to 1 and κ^c_0 equals $\frac{e^{wc_t}}{e^{wc_t}-1}$ when wc_t is evaluated at its long-run mean A_0.

D.2 Proof of Proposition 3

Setting Up Notation The starting point of the analysis is the Euler equation $E_t[M_{t+1}R^i_{t+1}] = 1$, where R^i_{t+1} denotes a gross return between dates t and $t+1$ on some asset i and M_{t+1} is the SDF. In logs:

$$E_t\left[m_{t+1}\right] + E_t\left[r^i_{t+1}\right] + \frac{1}{2}Var_t\left[m_{t+1}\right] + \frac{1}{2}Var_t\left[r^i_{t+1}\right] + Cov_t\left[m_{t+1}, r^i_{t+1}\right] = 0. \tag{50}$$

The same equation holds for the real risk-free rate $y_t(1)$, so that

$$y_t(1) = -E_t\left[m_{t+1}\right] - \frac{1}{2}Var_t\left[m_{t+1}\right]. \tag{51}$$

The expected excess return becomes:

$$E_t\left[r^{e,i}_{t+1}\right] = E_t\left[r^i_{t+1} - y_t(1)\right] + \frac{1}{2}Var_t\left[r^i_{t+1}\right] = -Cov_t\left[m_{t+1}, r^i_{t+1}\right] = -Cov_t\left[m_{t+1}, r^{e,i}_{t+1}\right], \tag{52}$$

where $r^{e,i}_{t+1}$ denotes the excess return on asset i corrected for the Jensen term.

We adopt the consumption growth specification of Bansal and Yaron (2004), repeated from the main text:

$$\Delta c_{t+1} = \mu_c + x_t + \sigma_t \eta_{t+1}, \tag{53}$$

$$x_{t+1} = \rho_x x_t + \varphi_e \sigma_t e_{t+1}, \tag{54}$$

$$\sigma^2_{t+1} = \bar{\sigma}^2 + \nu_1(\sigma^2_t - \sigma^2) + \sigma_w w_{t+1}, \tag{55}$$

where (η_t, e_t, w_t) are i.i.d. mean-zero, variance-one, normally distributed innovations.

Proof of Linearity The proof closely follows the proof in Bansal and Yaron (2004), henceforth BY, but adjusts all expressions for our timing of returns.

Proof. In what follows we focus on the return on a claim to aggregate consumption, denoted r^c, where

$$r^c_{t+1} = \kappa^c_0 + \Delta c_{t+1} + wc_{t+1} - \kappa^c_1 wc_t,$$

and derive the five terms in equation (50) for this asset.

Taking logs on both sides of the non-linear SDF expression in equation (46) of Appendix D.1 delivers an expression of the log SDF as a function of log consumption changes and the log total wealth return

$$m_{t+1} = \frac{1-\alpha}{1-\rho}\log\beta - \frac{1-\alpha}{1-\rho}\rho\Delta c_{t+1} + \left(\frac{1-\alpha}{1-\rho} - 1\right) r^c_{t+1}. \tag{56}$$

We conjecture that the log wealth-consumption ratio is linear in the two states x_t and $\sigma^2_t - \bar{\sigma}^2$,

$$wc_t = A^c_0 + A^c_1 x_t + A^c_2\left(\sigma^2_t - \bar{\sigma}^2\right).$$

As BY, we assume joint conditional normality of consumption growth, x, and the variance of consumption growth. We verify this conjecture from the Euler equation (50).

Using the conjecture for the wealth-consumption ratio, we compute innovations in the total wealth return, and its conditional mean and variance:

$$\begin{aligned}
r^c_{t+1} - E_t\left[r^c_{t+1}\right] &= \sigma_t\eta_{t+1} + A^c_1\varphi_e\sigma_t e_{t+1} + A^c_2\sigma_w w_{t+1} \\
E_t\left[r^c_{t+1}\right] &= r_0 + \left(1 - (\kappa^c_1 - \rho_x)A^c_1\right)x_t - A^c_2(\kappa^c_1 - \nu_1)\left(\sigma^2_t - \bar{\sigma}^2\right) \\
V_t\left[r^c_{t+1}\right] &= (1 + (A^c_1\varphi_e)^2)\sigma^2_t + (A^c_2)^2\sigma^2_w \\
r_0 &= \kappa^c_0 + A^c_0(1 - \kappa^c_1) + \mu_c
\end{aligned}$$

These equations correspond to (A.8) and (A.9) in the Appendix of BY.

Substituting in the expression for the log total wealth return r^c into the log SDF, we compute innovations, and the conditional mean and variance of the log SDF:

$$\begin{aligned}
m_{t+1} - E_t\left[m_{t+1}\right] &= -\alpha\sigma_t\eta_{t+1} - \frac{\alpha-\rho}{1-\rho}A_1^c\varphi_e\sigma_t e_{t+1} - \frac{\alpha-\rho}{1-\rho}A_2^c\sigma_w w_{t+1}, \\
E_t\left[m_{t+1}\right] &= m_0 - \rho x_t + \frac{\alpha-\rho}{1-\rho}(\kappa_1^c - \nu_1)A_2^c\left(\sigma_t^2 - \bar{\sigma}^2\right) \\
V_t\left[m_{t+1}\right] &= \left(\alpha^2 + \left(\frac{\alpha-\rho}{1-\rho}\right)^2 (A_1^c\varphi_e)^2\right)\sigma_t^2 + \left(\frac{\alpha-\rho}{1-\rho}\right)^2 (A_2^c)^2\sigma_w^2 \\
m_0 &= \frac{1-\alpha}{1-\rho}\log\beta - \frac{\alpha-\rho}{1-\rho}\left[\kappa_0^c + A_0^c(1-\kappa_1^c)\right] - \alpha\mu_c \qquad (57)
\end{aligned}$$

These expressions correspond to equations (A.10) and (A.27), and are only slightly different due to the different timing in returns.

The conditional covariance between the log consumption return and the log SDF is given by the conditional expectation of the product of their innovations

$$Cov_t\left[r_{t+1}^c, m_{t+1}\right] = E_t\left[r_{t+1}^c - E_t\left[r_{t+1}^c\right], m_{t+1} - E_t\left[m_{t+1}\right]\right] = \left(-\alpha - \frac{\alpha-\rho}{1-\rho}(A_1^c\varphi_e)^2\right)\sigma_t^2 - \frac{\alpha-\rho}{1-\rho}(A_2^c)^2\sigma_w^2$$

Using the method of undetermined coefficients and the five components of equation (50), we can solve for the constants A_0^c, A_1^c, and A_2^c:

$$\begin{aligned}
A_1^c &= \frac{1-\rho}{\kappa_1^c - \rho_x}, \qquad (58) \\
A_2^c &= \frac{(1-\rho)(1-\alpha)}{2(\kappa_1^c - \nu_1)}\left[1 + \frac{\varphi_e^2}{(\kappa_1^c - \rho_x)^2}\right], \qquad (59) \\
0 &= \frac{1-\alpha}{1-\rho}\left[\log\beta + \kappa_0^c + (1-\kappa_1^c)A_0^c\right] + (1-\alpha)\mu_c + \frac{1}{2}(1-\alpha)^2\left[1 + \frac{\varphi_e^2}{(\kappa_1^c - \rho_x)^2}\right]\bar{\sigma}^2 + \frac{1}{2}\left(\frac{1-\alpha}{1-\rho}\right)^2 (A_2^c)^2\sigma_w^2 \qquad (60)
\end{aligned}$$

The first two correspond to equations (A.5) and (A.7) in BY. The last equation implicitly solves A_0^c as a function of all parameters of the model. Because κ_0^c and κ_1^c are non-linear functions of A_0^c, this system of three equations needs to be solved simultaneously and numerically. Our computations indicate that the system has a unique solution. This verifies the conjecture that the log wealth-consumption ratio is linear in the two state variables. □

According to (52), the risk premium on the consumption claim is given by

$$E_t\left[r_{t+1}^e\right] = E_t\left[r_{t+1}^c - y_t(1)\right] + .5V_t[r_{t+1}^c] = -\lambda_{m,\eta}\sigma_t^2 + \lambda_{m,e}B\sigma_t^2 + \lambda_{m,w}A_2^c\sigma_w^2, \qquad (61)$$

This corresponds to equation (A.11) in BY, with $\lambda_{m,\eta} = -\alpha$, $\lambda_{m,e} = \frac{\alpha-\rho}{1-\rho}A_1^c\varphi_e$, and $\lambda_{m,w} = \frac{\alpha-\rho}{1-\rho}A_2^c$.

According to equation (51), the expression for the risk-free rate is given by

$$\begin{aligned}
y_t(1) &= h_0 + \rho x_t + h_1(\sigma_t^2 - \bar{\sigma}^2) \qquad (62) \\
h_0 &= -m_0 - .5\lambda_{m,w}^2\sigma_w^2 - .5\left(\lambda_{m,\eta}^2 + \lambda_{m,e}^2\right)\bar{\sigma}^2 \\
h_1 &= -\frac{\alpha-\rho}{1-\rho}(\kappa_1^c - \nu_1)A_2^c - .5\left(\lambda_{m,\eta}^2 + \lambda_{m,e}^2\right) \\
&= .5(\rho-\alpha)\left(1 + \frac{\varphi_e^2}{(\kappa_1^c - \rho_x)^2}\right) - .5\left(\alpha^2 + (\alpha-\rho)^2\frac{\varphi_e^2}{(\kappa_1^c - \rho_x)^2}\right)
\end{aligned}$$

This corresponds to equation (A.25-A.27) in BY. Its unconditional mean is simply h_0 (see A.28).

D.3 Quarterly Calibration LRR Model

The Bansal-Yaron model is calibrated and parameterized to monthly data. Since we want to use data on quarterly consumption and dividend growth, and a quarterly series for the wealth-consumption ratio, we recast the model at quarterly frequencies. We assume that the quarterly process for consumption growth, dividend growth, the low frequency component and the variance has the exact same structure than at the monthly frequency, with mean zero, standard deviation 1 innovations, but with different parameters. This appendix explains how the monthly parameters map into quarterly parameters. We denote all variables, shocks, and all parameters of the quarterly system with a tilde superscript. Our parameter values are listed at the end of this section, together with details on the simulation approach.

Preference Parameters Obviously, the preference parameters do not depend on the horizon ($\tilde{\alpha} = \alpha$ and $\tilde{\rho} = \rho$), except for the time discount factor $\tilde{\beta} = \beta^3$. Also, the long-run average log wealth-consumption ratio at the quarterly frequency is lower than at the monthly frequency by approximately log(3), because log of quarterly consumption is the log of three times monthly consumption.

Cash-flow Parameters We accomplish this by matching the conditional and unconditional mean and variance of log consumption and dividend growth. Log quarterly consumption growth is the sum of log consumption growth of three consecutive months. We obtain $\Delta\tilde{c}_{t+1} \equiv \Delta c_{t+3} + \Delta c_{t+2} + \Delta c_{t+1}$

$$\Delta\tilde{c}_{t+1} = 3\mu_c + (1+\rho_x+\rho_x^2)x_t + \sigma_t\eta_{t+1} + \sigma_{t+1}\eta_{t+2} + \sigma_{t+2}\eta_{t+3} + (1+\rho_x)\varphi_e\sigma_t e_{t+1} + \varphi_e\sigma_{t+1}e_{t+2} \tag{63}$$

Log quarterly dividend growth looks similar:

$$\Delta\tilde{d}_{t+1} = 3\mu_d + \phi(1+\rho_x+\rho_x^2)x_t + \varphi_d\sigma_t u_{t+1} + \varphi_d\sigma_{t+1}u_{t+2} + \varphi_d\sigma_{t+2}u_{t+3} + \phi(1+\rho_x)\varphi_e\sigma_t e_{t+1} + \phi\varphi_e\sigma_{t+1}e_{t+2} \tag{64}$$

First, we rescale the long-run component in the quarterly system, so that the coefficient on it in the consumption growth equation is still 1:

$$\tilde{x}_t = (1+\rho_x+\rho_x^2)x_t.$$

Second, we equate the unconditional mean of consumption and dividend growth :

$$\tilde{\mu} = 3\mu,\ \tilde{\mu}_d = 3\mu_d.$$

These imply that we also match the the conditional mean of consumption growth:

$$E_t[\Delta c_{t+3} + \Delta c_{t+2} + \Delta c_{t+1}] = 3\mu + (1+\rho_x+\rho_x^2)x_t = \tilde{\mu} + \tilde{x}_t = E_t[\Delta\tilde{c}_{t+1}]$$

Third, we also match the conditional mean of dividend growth by setting the quarterly leverage parameter $\tilde{\phi} = \phi$.

Fourth, we match the unconditional variance of quarterly consumption growth:

$$\begin{aligned}
V[\Delta\tilde{c}_{t+1}] &= (1+\rho_x+\rho_x^2)^2 V[x_t] + \sigma^2\left[3+(1+\rho_x)^2\varphi_e^2+\varphi_e^2\right] \\
&= (1+\rho_x+\rho_x^2)^2\frac{\varphi_e^2\sigma^2}{1-\rho_x^2} + \sigma^2\left[3+(1+\rho_x)^2\varphi_e^2+\varphi_e^2\right] \\
&= \frac{\tilde{\varphi}_e^2\tilde{\sigma}^2}{1-\tilde{\rho}_x^2} + \tilde{\sigma}^2
\end{aligned}$$

The first and second equalities use the law of iterated expectations to show that

$$V[\sigma_{t+j}\eta_{t+j+1}] \equiv E\left[E_{t+j}\left\{\sigma_{t+j}^2\eta_{t+j+1}^2\right\}\right] - \left(E\left[E_{t+j}\left\{\sigma_{t+j}\eta_{t+j+1}\right\}\right]\right)^2 = E\left[\sigma_{t+j}^2\right] - 0 = \sigma^2$$

and the same argument applies to terms of type $V[\sigma_{t+j}e_{t+j+1}]$. Coefficient matching on the variance of consumption expression delivers expressions for $\tilde{\sigma}^2$ and $\tilde{\varphi}_e$:

$$\tilde{\sigma}^2 = \sigma^2\left[3+(1+\rho_x)^2\varphi_e^2+\varphi_e^2\right]$$

$$\begin{aligned}
\tilde{\varphi}_e^2 &= \varphi_e^2\frac{(1-\tilde{\rho}_x^2)(1+\rho_x+\rho_x^2)^2}{1-\rho_x^2}\frac{\sigma^2}{\tilde{\sigma}^2} \\
&= \frac{(1-\rho_x^6)(1+\rho_x+\rho_x^2)^2}{1-\rho_x^2}\frac{\varphi_e^2}{3+(1+\rho_x)^2\varphi_e^2+\varphi_e^2},
\end{aligned}$$

where the third equality uses the first equality. Note that we imposed $\tilde{\rho}_x = \rho_x^3$, which follows from a desire to match the persistence of the long-run cash-flow component. Recursively substituting, we find that the three-month ahead x process has the following relationship to the current value:

$$x_{t+3} = \rho_x^3 x_t + \varphi_e\sigma_{t+2}e_{t+3} + \rho_x\varphi_e\sigma_{t+1}e_{t+2} + \rho_x^2\varphi_e\sigma_t e_{t+1}$$

which compares to the quarterly equation

$$\tilde{x}_{t+1} = \tilde{\rho}_x\tilde{x}_t + \tilde{\varphi}_e\tilde{\sigma}_t\tilde{e}_{t+1}$$

The two processes now have the same auto-correlation and unconditional variance.

Fifth, we match the unconditional variance of dividend growth. Given the assumptions we have made sofar, this pins down φ_d:

$$\tilde{\varphi}_d^2 = \frac{3\varphi_d^2 + \phi^2(1+\rho_x)^2\varphi_e^2 + \phi^2\varphi_e^2}{3+(1+\rho_x)^2\varphi_e^2+\varphi_e^2}$$

Sixth, we match the autocorrelation and the unconditional variance of economic uncertainty σ_t^2. Iterating forward, we obtain an expression that relates variance in month t to the one in month $t+3$:

$$\sigma_{t+3}^2 - \sigma^2 = \nu_1^3(\sigma_t^2-\sigma^2) + \sigma_w\nu_1^2 w_{t+1} + \sigma_w\nu_1 w_{t+2} + \sigma_w w_{t+3}$$

By setting $\tilde{\nu}_1 = \nu_1^3$ and $\tilde{\sigma}_w^2 = \sigma_w^2(1+\nu_1^2+\nu_1^4)$, we match the autocorrelation and variance of the quarterly equation

$$\tilde{\sigma}_{t+1}^2 - \tilde{\sigma}^2 = \tilde{\nu}_1(\tilde{\sigma}_t^2-\tilde{\sigma}^2) + \tilde{\sigma}_w\tilde{w}_{t+1}$$

Parameter Values We use $\rho = 2/3$, $\alpha = 10$, and $\beta = .997$ for preferences; and $\mu_c = .45e^{-2}$, $\bar{\sigma} = 1.35e^{-2}$, $\rho_x = .938$, $\varphi_e = .126$, $\nu_1 = .962$, and $\sigma_w = .39 * 10^{-5}$ for the cash-flow processes in (14)-(16). The vector $\Theta^{LRR} = (\alpha, \rho, \beta, \mu_c, \bar{\sigma}, \varphi_e, \rho_x, \nu_1, \sigma_w)$ stores these parameters. The corresponding monthly values are $\Theta^{LRR} = \left(10, .6666, .998985, .0015, .0078, .044, .979, .987, .23 * 10^{-5}\right)$. A simulation of the quarterly model recovered the annualized cash-flow and asset return moments of the monthly simulation.

The quarterly calibration implies the following solution to the system of equations in (58)-(60): $A_0^{c,LRR} = 5.85$, $A_1^{c,LRR} = [5.16, -175.10]$. The corresponding linearization constants are $\kappa_0^c = .0198$ and $\kappa_1^c = 1.0029$.

The quarterly parameters of the dividend claim are: $\phi = 3$ and $\varphi_d = 4.4960$.

Simulation Most population moments of interest are known in closed-form in the LRR model, so that we do not have to simulate. However, the simulation approach has the advantages of generating small-sample biases that may also exist in the data and delivering (bootstrap) standard errors.

We run 5,000 simulations of the model for 236 quarters each, corresponding to the period 1948-2006. In each simulation we draw a 236×3 matrix of mutually uncorrelated standard normal random variables for the cash-flow innovations (η, e, w) in (14)-(16). We start off each run at the steady-state ($x_0 = 0$ and $\sigma_t^2 = \bar{\sigma}^2$). For each run, we form log consumption growth Δc_t, the two state variables $[x_t, \sigma_t^2 - \bar{\sigma}^2]$, the log wealth-consumption ratio wc_t and its first difference, and the log total wealth return r_t^c. We compute their first and second moments. These moments are based on the last 220 quarters only, for consistency with the length for our data for consumption growth and the growth rate of the wealth-consumption ratio (1952.I-2006.IV). This has the added benefit that the first 16 quarters are "burn-in," so that the first observation we use for the state vector is different in each run.

D.4 Campbell-Shiller Decomposition

Expected discounted future returns and consumption growth rates are given by:

$$\begin{aligned} r_t^H &\equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j}\, r_{t+j}\right] = \frac{r_0^c}{\kappa_1^c - 1} + \frac{\rho}{\kappa_1^c - \rho_x} x_t - A_2^c(\sigma_t^2 - \bar{\sigma}^2) \\ \Delta c_t^H &\equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j}\, \Delta c_{t+j}\right] = \frac{\mu_c}{\kappa_1^c - 1} + \frac{1}{\kappa_1^c - \rho_x} x_t \end{aligned}$$

These expressions use the definitions of the total wealth return and consumption, as well as their dynamics. Starting from the affine expression for wc, we can write the wc ratio as a constant plus expected future consumption growth minus expected future total wealth returns:

$$\begin{aligned} wc_t &= A_0^c + A_1^c x_t + A_2^c(\sigma_t^2 - \bar{\sigma}^2) = A_0^c + \frac{1}{\kappa_1^c - \rho_x} x_t - \left(\frac{\rho}{\kappa_1^c - \rho_x} x_t - A_2^c(\sigma_t^2 - \bar{\sigma}^2)\right) \\ &= A_0^c + \left(\Delta c_t^H - \frac{\mu_c}{\kappa_1^c - 1}\right) - \left(r_t^H - \frac{r_0^c}{\kappa_1^c - 1}\right) = \frac{\kappa_0^c}{\kappa_1^c - 1} + \Delta c_t^H - r_t^H \end{aligned}$$

The second equality uses the definition of A_1^c. The third equality uses the definition of r_t^H and Δc_t^H. The fourth equality uses the definition of r_0^c.

The variance of the log wealth-consumption ratio can be written in two equivalent ways:

$$V\left[\Delta c_t^H\right] + V\left[r_t^H\right] - 2Cov\left[r_t^H, \Delta c_t^H\right] = V\left[wc_t\right] = Cov\left[wc_t, \Delta c_t^H\right] + Cov\left[wc_t, -r_t^H\right]$$

In the LRR model, the terms in this expression are given by

$$
\begin{aligned}
V\left[\Delta c_t^H\right] &= \frac{1}{(\kappa_1^c-\rho_x)^2}\frac{\varphi_e^2}{1-\rho_x^2}\bar{\sigma}^2 > 0 \\
V\left[r_t^H\right] &= \frac{\rho^2}{(\kappa_1^c-\rho_x)^2}\frac{\varphi_e^2}{1-\rho_x^2}\bar{\sigma}^2 + (A_2^c)^2\frac{\sigma_w^2}{1-\nu_1^2} > 0 \\
Cov\left[r_t^H, \Delta c_t^H\right] &= \frac{\rho}{(\kappa_1^c-\rho_x)^2}\frac{\varphi_e^2}{1-\rho_x^2}\bar{\sigma}^2 > 0 \\
Cov\left[wc_t, \Delta c_t^H\right] &= \frac{1-\rho}{(\kappa_1^c-\rho_x)^2}\frac{\varphi_e^2}{1-\rho_x^2}\bar{\sigma}^2 > 0 \quad \Leftrightarrow \rho < 1 \\
Cov\left[wc_t, -r_t^H\right] &= \frac{\rho^2-\rho}{(\kappa_1^c-\rho_x)^2}\frac{\varphi_e^2}{1-\rho_x^2}\bar{\sigma}^2 + (A_2^c)^2\frac{\sigma_w^2}{1-\nu_1^2} > 0 \quad \Leftarrow \rho > 1
\end{aligned}
$$

We can break up expected future returns into a risk-free rate component and a risk premium component. The former is equal to

$$
E_t\left[\sum_{j=1}^{\infty}(\kappa_1^c)^{-j}\, r_{t+j-1}^f\right] = \frac{h_0}{\kappa_1^c-1} + \frac{\rho}{\kappa_1^c-\rho_x}x_t + \frac{h_1}{\kappa_1^c-\nu_1}(\sigma_t^2-\bar{\sigma}^2), \tag{65}
$$

where the second equation uses the expression for the risk-free rate in equation (62) to compute future risk-free rates and takes their time-t expectations. The risk premium component is simply the difference between r_t^H and the second expression.

D.5 Pricing Stocks in the LRR Model

Dividend Growth Process We start by pricing a claim to aggregate dividends, where the dividend process follows the specification in Bansal and Yaron (2004):

$$
\Delta d_{t+1} = \mu_d + \phi x_t + \varphi_d \sigma_t u_{t+1} \tag{66}
$$

The shock u_t is orthogonal to (η, e, w).

Defining returns ex-dividend and using the Campbell (1991) linearization, the log return on a claim to the aggregate dividend can be written as:

$$
r_{t+1}^m = \Delta d_{t+1} + pd_{t+1}^m + \kappa_0^m - \kappa_1^m pd_t^m,
$$

with coefficients

$$
\kappa_1^m = \frac{e^{A_0^m}}{e^{A_0^m}-1} > 1, \quad \text{and} \quad \kappa_0^m = -\log\left(e^{A_0^m}-1\right) + \frac{e^{A_0^m}}{e^{A_0^m}-1}A_0^m
$$

which depend on the long-run log price-dividend ratio A_0^m.

Proof of Linearity We conjecture, as we did for the wealth-consumption ratio, that the log price dividend ratio is linear in the two state variables:

$$
pd_t^m = A_0^m + A_1^m x_t + A_2^m(\sigma_t^2 - \bar{\sigma}^2).
$$

As we did for the return on the consumption claim, we compute innovations in the dividend claim return, and its conditional mean and variance:

$$
\begin{aligned}
r_{t+1}^m - E_t\left[r_{t+1}^m\right] &= \varphi_d \sigma_t u_{t+1} + \beta_{m,e}\sigma_t e_{t+1} + \beta_{m,w}\sigma_w w_{t+1} \\
E_t\left[r_{t+1}^m\right] &= r_0^m + [\phi + A_1^m(\rho_x - \kappa_1^m)]x_t - A_2^m(\kappa_1^m - \nu_1)\left(\sigma_t^2 - \bar{\sigma}^2\right) \\
V_t\left[r_{t+1}^m\right] &= (\varphi_d^2 + \beta_{m,e}^2)\sigma_t^2 + \beta_{m,w}^2\sigma_w^2, \\
r_0^m &= \kappa_0^m + A_0^m(1-\kappa_1^m) + \mu_d
\end{aligned}
$$

where $\beta_{m,e} = A_1^m \varphi_e$ and $\beta_{m,w} = A_2^m$. These equations correspond to (A.12) and (A.13) in the Appendix of Bansal and Yaron (2004). Finally, the conditional covariance between the log SDF and the log dividend claim return is

$$
Cov_t\left[m_{t+1}, r_{t+1}^m\right] = -\lambda_{m,e}\beta_{m,e}\sigma_t^2 - \lambda_{m,w}\beta_{m,w}\sigma_w^2
$$

From the Euler equation for this return $E_t\left[m_{t+1}\right] + E_t\left[r_{t+1}^m\right] + \frac{1}{2}V_t\left[m_{t+1}\right] + \frac{1}{2}V_t\left[r_{t+1}^m\right] + Cov_t[m_{t+1}, r_{t+1}^m] = 0$ and the method of undetermined coefficients, we can use the same procedure as described in D.2, and solve for the constants A_0^m, A_1^m, and A_2^m:

$$
\begin{aligned}
A_1^m &= \frac{\phi - \rho}{\kappa_1^m - \rho_x}, \\
A_2^m &= \frac{\left[\frac{\alpha-\rho}{1-\rho}A_2^c(\kappa_1 - \nu_1) + .5H_m\right]}{\kappa_1^m - \nu_1}, \\
0 &= m_0 + \kappa_0^m + (1-\kappa_1^m)A_0^m + \mu_d + \frac{1}{2}H_m\bar{\sigma}^2 + \frac{1}{2}\left(A_2^m - A_2^c\frac{\alpha-\rho}{1-\rho}\right)^2\sigma_w^2
\end{aligned}
$$

where

$$
H_m = \alpha^2 + \left(A_1^m\varphi_e + (\frac{\rho-\alpha}{1-\rho})A_1^c\varphi_e\right)^2 + \varphi_d^2
$$

Again, this is a non-linear system in three equations and three unknowns, which we solve numerically. The first two equations correspond to (A.16) and (A.20) in BY.

Equity Risk premium and CS Decomposition The equity risk premium on the dividend claim (adjusted for a Jensen term) becomes:

$$
E_t\left[r_{t+1}^{e,m}\right] \equiv E_t\left[r_{t+1}^m - y_t(1)\right] + .5V_t[r_{t+1}^m] = \lambda_{m,e}\beta_{m,e}\sigma_t^2 + \lambda_{m,w}\beta_{m,w}\sigma_w^2 \tag{67}
$$

This corresponds to equation (A.14) in BY.

Expected discounted future equity returns and dividend growth rates are given by:

$$
r_t^{m,H} \equiv E_t\left[\sum_{j=1}^{\infty}(\kappa_1^m)^{-j} r_{t+j}^m\right] = \frac{r_0^m}{\kappa_1^m - 1} + \frac{\rho}{\kappa_1^m - \rho_x}x_t - A_2^m(\sigma_t^2 - \bar{\sigma}^2) \tag{68}
$$

$$
\Delta d_t^H \equiv E_t\left[\sum_{j=1}^{\infty}(\kappa_1^m)^{-j}\Delta d_{t+j}\right] = \frac{\mu_d}{\kappa_1^m - 1} + \frac{\phi}{\kappa_1^m - \rho_x}x_t \tag{69}
$$

From these expressions, it is easy to see that $pd_t^m = \frac{\kappa_0^m}{\kappa_1^m - 1} + \Delta d_t^H - r_t^{m,H}$, and to compute the elements of the variance-decomposition:

$$V[pd_t^m] = Cov[pd_t^m, \Delta d_t^H] + Cov[pd_t^m, -r_t^{m,H}] = V[\Delta d_t^H] + V[r_t^{m,H}] - 2Cov[\Delta d_t^H, r_t^{m,H}].$$

E The External Habit Model

The organization of this EH model appendix exactly parallels the treatment of the LRR model in Appendix (D.2).

E.1 Proof of Proposition 4

Proof. We conjecture that the log wealth-consumption ratio is linear in the sole state variable $(s_t - \bar{s})$,

$$wc_t = A_0^c + A_1^c (s_t - \bar{s}).$$

As Campbell and Cochrane (1999), henceforth CC, we assume joint conditional normality of consumption growth and the surplus consumption ratio. We verify this conjecture from the Euler equation for total wealth.
We start from the canonical log SDF in the external habit model:

$$m_{t+1} = \log\beta - \alpha\Delta c_{t+1} - \alpha\Delta s_{t+1}.$$

Substituting in the expression for returns into the log SDF, we compute innovations, and the conditional mean and variance of the log SDF:

$$\begin{aligned} m_{t+1} - E_t[m_{t+1}] &= -\alpha(1+\lambda_t)\bar{\sigma}\eta_{t+1} \\ E_t[m_{t+1}] &= m_0 + \alpha(1-\rho_s)(s_t - \bar{s}) \\ V_t[m_{t+1}] &= \alpha^2 (1+\lambda_t)^2 \bar{\sigma}^2 \\ m_0 &= \log\beta - \alpha\mu_c \end{aligned} \tag{70}$$

Likewise, we compute innovations in the consumption claim return, and its conditional mean and variance:

$$\begin{aligned} r_{t+1}^c - E_t\left[r_{t+1}^c\right] &= (1 + A_1^c\lambda_t)\bar{\sigma}\eta_{t+1} \\ E_t\left[r_{t+1}^c\right] &= r_0 - A_1^c(\kappa_1^c - \rho_s)(s_t - \bar{s}) \\ V_t\left[r_{t+1}^c\right] &= (1 + A_1^c\lambda_t)^2\bar{\sigma}^2 \\ r_0 &= \kappa_0^c + A_0^c(1-\kappa_1^c) + \mu_c \end{aligned} \tag{71}$$

The conditional covariance between the log consumption return and the log SDF is given by the conditional expectation of the product of their innovations

$$Cov_t\left[m_{t+1}, r_{t+1}^c\right] = -\alpha(1+\lambda_t)(1 + A_1^c\lambda_t)\bar{\sigma}^2$$

We assume that the sensitivity function takes the following form

$$\lambda_t = \frac{\bar{S}^{-1}\sqrt{1 - 2(s_t - \bar{s})} + 1 - \alpha}{\alpha - A_1^c}$$

Using the method of undetermined coefficients and the five components of equation (50), we can solve for the constants A_0^c and A_1^c:

$$A_1^c = \frac{(1-\rho_s)\alpha - \bar{\sigma}^2 \bar{S}^{-2}}{\kappa_1^c - \rho_s}, \tag{72}$$

$$0 = \log\beta + \kappa_0^c + (1-\kappa_1^c)A_0^c + (1-\alpha)\mu_c + .5\bar{\sigma}^2\bar{S}^{-2} \tag{73}$$

This verifies that our conjecture was correct. □

It follows immediately that the log SDF can be written as a function of consumption growth and the change in the log wealth-consumption ratio

$$m_{t+1} = \log\beta - \alpha\Delta c_{t+1} - \frac{\alpha}{A_1^c}\left(wc_{t+1} - wc_t\right).$$

The risk premium on the consumption claim is given by $Cov_t[r_{t+1}^c, -m_{t+1}]$:

$$E_t\left[r_{t+1}^e\right] \equiv E_t\left[r_{t+1}^c - y_t(1)\right] + .5V_t[r_{t+1}^c] = \alpha\left(1+\lambda_t\right)\left(1 + A_1^c\lambda_t\right)\bar{\sigma}^2, \tag{74}$$

where the second term on the left is a Jensen adjustment. The expression for the risk-free rate appears in the next section E.2.

E.2 The Steady-State Habit Level

Campbell and Cochrane (1999) engineer their sensitivity function λ_t to deliver a risk-free rate that is linear in the state $s_t - \bar{s}$. (They mostly study a special case with a constant risk-free rate.) The linearity of the risk-free rate is accomplished by choosing

$$\lambda_t^{CC} = \bar{S}^{-1}\sqrt{1 - 2\left(s_t - \bar{s}\right)} - 1 \tag{75}$$

Note that if the risk aversion parameter $\alpha = 2$ and $A_1^c = 1$, our sensitivity function exactly coincides with CC's. Instead, we engineer our sensitivity function to deliver a log wealth-consumption ratio that is linear in $s_t - \bar{s}$.
As a result of our choice, the risk-free rate, $y_t(1) = -E_t\left[m_{t+1}\right] - .5V_t\left[m_{t+1}\right]$, is no longer linear in the state, but contains an additional square-root term:

$$y_t(1) = h_0 + \left[\frac{\bar{\sigma}^2\alpha^2\bar{S}^{-2}}{(\alpha - A_1^c)^2} - \alpha(1-\rho_s)\right](s_t - \bar{s}) - \bar{\sigma}^2\alpha^2\frac{(1-A_1^c)\bar{S}^{-1}}{(\alpha - A_1^c)^2}\left(\sqrt{1 - 2\left(s_t - \bar{s}\right)} - 1\right) \tag{76}$$

$$h_0 = -\log\beta + \alpha\mu_c - .5\bar{\sigma}^2\alpha^2(1+\lambda(\bar{s}))^2, \quad \text{where} \quad \lambda(\bar{s}) = \left(\frac{\bar{S}^{-1} + 1 - \alpha}{\alpha - A_1^c}\right) \tag{77}$$

where $\lambda(\bar{s})$ is obtained from evaluating our sensitivity function at $s_t = \bar{s}$.
CC obtain a similar expression, but without the last term. If $\alpha = 2$ and $A_1^c = 1$, the expression collapses to the one in CC. A constant risk-free rate obtains in the CC model when $\bar{S}^{-1} = \bar{\sigma}^{-1}\sqrt{\frac{1-\rho_s}{\alpha}}$ because this choice makes the linear term vanish. While there is no $\bar{S}$ that guarantees a constant risk-less interest rate under our assumptions, we choose $\bar{S}$ to match the *steady-state* risk-free rate in CC, $\bar{r} = -\log\beta + \alpha\mu - .5\alpha(1-\rho_s)$. That is, we set $s_t = \bar{s}$ in the above equation, which then collapses to h_0. Setting $\bar{r} = h_0$ allows us to solve for $\bar{S}^{-1}$ as a function of A_1^c and the structural parameters α, ρ_s, and $\bar{\sigma}$:

$$\bar{S}^{-1} = (\alpha - A_1^c)\left(\bar{\sigma}^{-1}\sqrt{\frac{1-\rho_s}{\alpha}}\right) - 1 + A_1^c. \tag{78}$$

Substituting this expression back into the sensitivity function (20), we find that the steady-state sensitivity level $\lambda(\bar{s}) = \bar{\sigma}^{-1}\sqrt{\frac{1-\rho_s}{\alpha}} - 1$. This implies that we generate the same steady-state conditional covariance between the surplus consumption ratio and consumption growth as in CC.

As in CC, we define a maximum value for the log surplus consumption ratio s_{max}, as the value at which λ_t runs into zero:

$$s_{max} = \bar{s} + \frac{1}{2}\left(1 - (\alpha-1)^2\bar{S}^2\right)$$

Note that if $\alpha = 2$, this coincides with equation (11) in CC. It is understood that $\lambda_t = 0$ for $s_t \geq s_{max}$.

E.3 Alternative Way of Pinning Down $\bar{S}$

To conclude the discussion of the EH model, we investigate an alternative way to pin down $\bar{S}$. In our benchmark method, outlined in Appendix E.2, we chose it to match the steady state risk-free rate in Campbell and Cochrane (1999). Here, the alternative is to pin down $\bar{S}$ to match the average wealth-consumption ratio of 26.75 in Campbell and Cochrane (1999). As before, we solve a system of three equations in $(A_0^c, A_1^c, \bar{S})$, only the third of which is different and simply imposes that $e^{A_0^c - \log(4)} = 26.75$. We obtain the following solution: $A_0^c = 4.673$, $A_1^c = 0.447$, and $\bar{S} = 0.0339$. The wealth-consumption ratio is higher and less sensitive to the surplus-consumption ratio than in the benchmark case. The volatility of the surplus-consumption ratio is 41.6%, similar to the benchmark model. Because A_1^c is lower, so is the volatility of the wc ratio. It is 18.6% in the model, still higher than in the data. The volatilities of the change in the wc ratio and of the total wealth return are also lowered, but remain too high. The consumption risk premium is down from 2.67% per quarter to 1.97% per quarter, much above those in the data. This alternative calibration has two drawbacks. First, the risk-free rate turns negative to -1.2% per year. This can happen because we are no longer pinning $\bar{S}$ down to match the steady-state risk-free rate. The risk-free rate is also more volatile: .59% versus .03 in the main text, but actually closer to the .55% in the data. The second drawback of this calibration is that the price-dividend ratio on equity is too smooth. The volatility of pd^m is now only 12.5% per quarter compared to 27% in the data. (The equity premium goes down from 3.30% per quarter to 2.23%.)

E.4 Quarterly Calibration EH Model

Preference Parameters Again, the preference parameter does not depend on the horizon ($\tilde{\alpha} = \alpha$, except for the time discount factor $\tilde{\beta} = \beta^3$). The surplus consumption ratio has the same law of motion as in the monthly model, but we set its persistence equal to $\tilde{\rho}_s = \rho_s^3$.

Cash-flow Parameters Following a similar logic, we can match mean and variance of quarterly consumption and dividend growth in the CC model. From matching the means we get:

$$\tilde{\mu} = 3\mu,\ \tilde{\mu}_d = 3\mu_d.$$

From matching the variances we get

$$\tilde{\sigma}^2 = 3\sigma^2,\ \tilde{\varphi}_d = \varphi_d,\ \tilde{\chi} = \chi.$$

Calibration We work with a quarterly calibration and set $\alpha = 2$, $\rho_s = .9658$, and $\beta = .971$ for preferences, and $\mu_c = .47e^{-2}$ and $\bar{\sigma} = .75e^{-2}$ for the cash-flow process (18). We summarize the parameters in the vector $\Theta^{EH} = (\alpha, \rho_s, \beta, \mu_c, \bar{\sigma})$. The corresponding monthly values are $\Theta^{EH} = (\alpha, \rho_s, \beta, \mu_c, \bar{\sigma}) = \left(2, .9885, .990336, .1575e^{-2}, .433e^{-2}\right)$.

We solve for the loadings of the state variables in the log wealth-consumption ratio form equations (72), (73), and (78), and find: $A_0^{EH} = 3.86$, $A_1^{EH} = 0.778$, and $\bar{S} = .0474$. The corresponding Campbell-Shiller linearization constants are $\kappa_0^c = .1046$ and $\kappa_1^c = 1.021583$.

For the dividend process described below, we set $\mu_d = \mu$, $\varphi_d = 7.32$, and $\chi = 0.20$.

A simulation of the quarterly model recovered the annualized cash-flow and asset return moments of the monthly simulation.

E.5 Campbell-Shiller Decomposition

Using (71) and the law of motion for s_t and consumption growth, expected discounted future returns and consumption growth rates are given by:

$$
\begin{aligned}
r_t^H &\equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j}\, r_{t+j}\right] = \frac{r_0^c}{\kappa_1^c - 1} - A_1^c(s_t - \bar{s}) \\
\Delta c_t^H &\equiv E_t\left[\sum_{j=1}^{\infty} (\kappa_1^c)^{-j}\, \Delta c_{t+j}\right] = \frac{\mu_c}{\kappa_1^c - 1}
\end{aligned}
$$

These expressions enable us to go back and forth between the affine log wealth-consumption ratio expression and the Campbell-Shiller decomposition:

$$
\begin{aligned}
wc_t &= A_0^c + A_1^c(s_t - \bar{s}) = A_0^c + \left(\Delta c_t^H - \frac{\mu_c}{\kappa_1^c - 1}\right) - \left(r_t^H - \frac{r_0^c}{\kappa_1^c - 1}\right) \\
&= \frac{\kappa_0^c}{\kappa_1^c - 1} + \Delta c_t^H - r_t^H,
\end{aligned}
$$

The variance of the log wealth-consumption ratio can be written in two equivalent ways:

$$
V\left[\Delta c_t^H\right] + V\left[r_t^H\right] - 2Cov\left[r_t^H, \Delta c_t^H\right] = V\left[wc_t\right] = Cov\left[wc_t, \Delta c_t^H\right] + Cov\left[wc_t, -r_t^H\right]
$$

In the EH model, the terms in this expression are given by

$$
\begin{aligned}
V\left[\Delta c_t^H\right] &= 0,\ \ Cov\left[r_t^H, \Delta c_t^H\right] = 0,\ \ Cov\left[wc_t, \Delta c_t^H\right] = 0 \\
V\left[r_t^H\right] &= (A_1^c)^2 \left(\frac{\bar{S}^{-1} + 1 - \alpha}{\alpha - A_1^c}\right)^2 \frac{1}{1 - \rho_s^2}\bar{\sigma}^2 > 0 \\
Cov\left[wc_t, -r_t^H\right] &= (A_1^c)^2 \left(\frac{\bar{S}^{-1} + 1 - \alpha}{\alpha - A_1^c}\right)^2 \frac{1}{1 - \rho_s^2}\bar{\sigma}^2 > 0
\end{aligned}
$$

Likewise, there is no predictability in dividend growth (see equation 79). Therefore, $V[pd_t] = V\left[r_t^{H,m}\right]$, where the latter is the unconditional variance of the expected return on the dividend claim.

E.6 Pricing Stocks in EH Model

The main difference with the analysis in the long-run risk model, and the analysis for the total wealth return in the EH model is that the return to the aggregate dividend claim *cannot* be written as a linear function of the state variables. Our choice of the sensitivity function makes the log wealth-consumption ratio linear in the surplus consumption ratio. But, for that same sensitivity function, the log price-dividend ratio is not linear in the surplus-consumption ratio. As a result, we need to resort to a non-linear computation of the price-dividend ratio on the aggregate dividend claim.

Dividend Growth Process In Campbell and Cochrane (1999), dividend growth is i.i.d., with the same mean μ as consumption growth, and innovations that are correlated with the innovations in consumption growth. To make the dividend growth process more directly comparable across models, we write it as a function of innovations to consumption growth η and innovations u that are orthogonal to η:

$$\Delta d_{t+1} = \mu_d + \varphi_d \bar{\sigma} u_{t+1} + \varphi_d \bar{\sigma} \chi \eta_{t+1}. \tag{79}$$

It follows immediately that its (un)conditional variance equals $\varphi_d^2 \bar{\sigma}^2 (1 + \chi^2)$ and its (un)conditional covariance with consumption growth is $\varphi_d \bar{\sigma}^2 \chi$. If correlation between consumption and dividend growth is $corr$, then $\chi = \sqrt{corr^2/(1 - corr^2)}$. We set φ_d and χ to replicate the unconditional variance of dividend growth and the correlation of dividend growth and consumption growth $corr$ in Campbell and Cochrane (1999).

Computation of Price-Dividend Ratio Wachter (2005) shows that the price-dividend ratio on a claim to aggregate dividends can be written as the sum of the price-dividend ratios on strips to the period-n dividend, for $n = 1, \cdots, \infty$:

$$\frac{P_t}{D_t} = \sum_{n=1}^{\infty} \frac{P_{nt}^d}{D_t} \tag{80}$$

We adopt her methodology and show it continues to hold for our slightly different dividend growth process in equation (79). A similar approach works when there is cointegration between consumption and dividends (Appendix A in Wachter (2005)).

The Euler equation for the period-n strip delivers the following expression for the price-dividend ratio

$$\frac{P_{nt}^d}{D_t} = E_t \left[M_{t+1} \frac{P_{n-1,t+1}^d}{D_{t+1}} \frac{D_{t+1}}{D_t} \right]$$

We conjecture that the price-dividend ratio on the period-n strip equals a function $F_n(s_t)$, which follows the recursion

$$F_n(s_t) = \beta e^{\mu_d - \alpha\mu_c + \alpha(1-\rho_s)(s_t - \bar{s}) + \frac{1}{2}\varphi_d^2 \bar{\sigma}^2} E_t \left[e^{[\varphi_d \chi - \alpha(1+\lambda_t)]\bar{\sigma}\eta_{t+1}} F_{n-1}(s_{t+1}) \right],$$

starting at $F_0(s_t) = 1$. We now verify this conjecture.

Proof. Substituting in the conjecture $\frac{P_{nt}^d}{D_t} = F_n(s_t)$ into the Euler equation for the period-n strip, we get

$$F_n(s_t) = E_t \left[M_{t+1} F_{n-1}(s_{t+1}) \frac{D_{t+1}}{D_t} \right].$$

Substituting in for the stochastic discount factor M and the dividend growth process (79), this becomes

$$F_n(s_t) = \beta e^{\mu_d - \alpha\mu_c + \alpha(1-\rho_s)(s_t - \bar{s})} E_t \left[e^{-\alpha(1+\lambda_t)\bar{\sigma}\eta_{t+1}} F_{n-1}(s_{t+1}) e^{\varphi_d \bar{\sigma} u_{t+1} + \varphi_d \chi \bar{\sigma} \eta_{t+1}} \right].$$

Because u and η are independent, we can write the expectation as a product of expectations. Because u is standard normal, the expectation in the previous expression can be written as

$$e^{\frac{1}{2}\varphi_d^2\bar{\sigma}^2} E_t \left[e^{[\varphi_d \chi - \alpha(1+\lambda_t)]\bar{\sigma}\eta_{t+1}} F_{n-1}(s_{t+1}) \right].$$

This then verifies the conjecture for $F_n(s_t)$. $\square$

Finally, let $g(\eta)$ be the standard normal pdf, then we can compute this function through numerical integration

$$F_n(s_t) = \beta e^{\mu_d - \alpha\mu_c + \alpha(1-\rho_s)(s_t - \bar{s}) + \frac{1}{2}\varphi_d^2\bar{\sigma}^2} \int_{-\infty}^{+\infty} e^{[\varphi_d \chi - \alpha(1+\lambda(s_t))]\bar{\sigma}\eta_{t+1}} F_{n-1}(s_{t+1}) g(\eta_{t+1}) d\eta_{t+1},$$

starting at $F_0(s_t) = 1$. The grid for s_t includes 14 very low values for s_t (-300, -250, -200, -150, -100, -50, -40, -30, -20, -15, -10, -9, -8, -7), 100 linearly spaced points between -6.5 and $\bar{s} * 1.001 = -2.85$, and the log of 100 linearly spaced points between $\bar{S}$ and $\exp(1.0000001 s_{max})$. The function evaluation $F_{n-1}(s_{t+1})$ is done using linear interpolation (and extrapolation) on the grid for the log surplus-consumption ratio s. The integral is computed in matlab using quad.m. The price dividend ratio is computed as the sum of the price-dividend ratios for the first 500 strips.