

NBER WORKING PAPER SERIES

THE WEALTH-CONSUMPTION RATIO

Hanno Lustig
Stijn Van Nieuwerburgh
Adrien Verdelhan

Working Paper 13896
<http://www.nber.org/papers/w13896>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2008

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, 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 and conference participants at the SED in Prague, the CEPR meeting in Gerzensee, the EFA meeting in Ljubljana, and the AFA and AEA meetings in New Orleans for comments. This work is supported by the National Science Foundation under Grant No 0550910. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2008 by Hanno Lustig, Stijn Van Nieuwerburgh, and Adrien Verdelhan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Wealth-Consumption Ratio

Hanno Lustig, Stijn Van Nieuwerburgh, and Adrien Verdelhan

NBER Working Paper No. 13896

March 2008, Revised May 2008

JEL No. E21,G10,G12

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 overall, but most of the return predictability is for future interest rates, not excess returns. We conclude that the properties of the total wealth portfolio are more similar to those of a long-maturity bond portfolio than those of a stock portfolio. The differences that we find between the risk-return characteristics of equity and total wealth suggest that equity is a special asset class.

Hanno Lustig
Department of Economics
UCLA
Box 951477
Los Angeles, CA 90095-1477
and NBER
hlustig@econ.ucla.edu

Adrien Verdelhan
Department of Economics,
Boston University,
270 Bay State Road,
Boston, MA 02215
av@bu.edu

Stijn Van Nieuwerburgh
Stern School of Business
New York University
44 W 4th Street, Suite 9-120
New York, NY 10012
and NBER
svnieuwe@stern.nyu.edu

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 the various types 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% per year, compared to a much higher equity risk premium of 6.9%. 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 (wc) ratio is less volatile than the price-dividend ratio: its standard deviation is 17% versus 27%. The return on total wealth has a volatility that is 9.8% per year, compared to 16.7% for equity returns. Our estimation produces a variance decomposition of the wc 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 wc 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 total wealth returns is variation in future risk-free rates, and not variation in future excess returns. In contrast, the price-dividend ratio on equity mostly predicts future excess equity returns.

This gap between the properties of total wealth and equity is crucial for the evaluation of dynamic asset pricing theories. In the Capital Asset Pricing Model, the total wealth return is the right pricing factor (Roll 1977). In the Inter-temporal CAPM, current and future total wealth returns can substitute for consumption growth as pricing factors (Campbell 1993). However, applied work commonly tests dynamic asset pricing models (DAPMs) by using the stock market return as a proxy for the total wealth return. This is problematic because the stock market return turns out to be a poor proxy for the total wealth return. 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. 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), and a no-arbitrage vector auto-regression (VAR) to describe the dynamics of stock returns, bond yields, and consumption growth.

While we observe the cash flow on human wealth (labor income), we do not observe its discount rate (expected return). Therefore, its price is unknown. For housing wealth and other parts of

broad financial wealth such as private business wealth, there is a lack of reliable market price data. Our approach avoids making somewhat arbitrary assumptions about expected returns to value these holdings. Instead, we infer the conditional market prices of different types of aggregate risk from stock and bond prices. Armed with these estimates and with an empirical model for the dynamics of aggregate consumption growth, we value a claim to aggregate consumption.

Our work embeds the VAR methodology of Campbell (1991, 1993, 1996) into the no-arbitrage framework of Ang and Piazzesi (2003). Like Campbell (1993), we specify the state variables that are in the investor's information set and we assume that their dynamics are given by a VAR system. 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 risk. The market prices of risk 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) and Dai and Singleton (2000). 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. In sum, our model provides a close fit for the risk premia on bonds and stocks. With the prices of risk inferred from traded assets, we price a claim to aggregate consumption and aggregate labor income.

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.

In the benchmark model, we do 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 Alvarez and Jermann (2004), 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 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. We note that 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 1999.

On average, the risk-return properties of human wealth closely resemble those of total wealth. We estimate human wealth to be 90% of total wealth. This estimate is consistent with Jorgenson and Fraumeni (1989), whose calculations also suggest a 90% 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 labor income, the length of a typical career, 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 career 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 two leading DAPMs. Our goal is not to try to statistically test these models on the basis of our wc estimates, but simply to highlight the key role of the wc ratio in leading asset pricing models. Interestingly, the external habit model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and Yaron (2004) have very different predictions for the wealth-consumption ratio despite their similar predictions for equity 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

¹Early 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. Chen and Ludvigson (2007) estimate the habit process for a class of EH models.

²Hansen, Heaton, and Li (2008), Parker and Julliard (2005) and Malloy (2005) 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, Goldstein, and Collin-Dufresne (2005) for credit spreads, and Colacito and Croce (2005) and Bansal and Shaliastovich (2007) for international finance. Martin (2007) allows for higher order moments in consumption growth. 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.

volatile wealth-consumption ratio than the price-dividend ratio on equity. The average wealth-consumption ratio in the benchmark LRR model is 87. This is the exact same value we estimate in the data; it shows that our numbers are consistent with a standard general equilibrium asset pricing model. On the predictability side, the LRR model delivers less predictability in total wealth returns than in equity returns, and most of the return predictability comes through the risk-free rate. These properties are also consistent with the data. However, the LRR model generates too much consumption growth predictability. On the other hand, 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. On the predictability side, the variance decomposition of the wc ratio replicates the low consumption growth predictability of the data. The EH model generates substantial time variation in expected equity returns. Because of the similarity of the consumption and dividend claims, this translates into substantial total wealth (excess) return predictability. Our results suggest that DAPMs should aim to generate different properties for total wealth, the price of a claim to consumption, and equity, the price of a claim to dividends.

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 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 wc ratio. Also, it uses the aggregate household wealth data 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 (2007) 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 business cycles from asset prices. Their log SDF is linear in aggregate consumption growth and the market return, and their model is calibrated to match only the unconditional equity premium. Their model does not allow for time-varying risk premia. They estimate a much smaller consumption risk premium of 0.2%, 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% and lowers the wealth-consumption ratio substantially.

We start the paper 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:

$$W_{t+1} = R_{t+1}^c (W_t - C_t). \quad (1)$$

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_{t+1}^c 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

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

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 $r_t^c = \log(R_t^c)$ around the long-run average log wealth-consumption ratio $A_0^c \equiv E[w_t - c_t]$.³

$$r_{t+1}^c = \kappa_0^c + \Delta c_{t+1} + wc_{t+1} - \kappa_1^c wc_t, \quad (2)$$

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

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

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

The linearization constants κ_0^c and κ_1^c are non-linear functions of the unconditional mean wealth-consumption ratio A_0^c :

$$\kappa_1^c = \frac{e^{A_0^c}}{e^{A_0^c} - 1} > 1 \quad \text{and} \quad \kappa_0^c = -\log(e^{A_0^c} - 1) + \frac{e^{A_0^c}}{e^{A_0^c} - 1} A_0^c. \quad (3)$$

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.

Let aggregate consumption $C_t(z^t)$ and aggregate labor income $L_t(z^t)$ depend on the history of the aggregate shocks $z \in Z$, $z^t = \{z_0, z_1, \dots, z_t\}$. 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 assume that the traded asset payoffs span the aggregate shocks, but not the idiosyncratic shocks. 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:

$$\frac{\Lambda_t^*}{\Lambda_{t-1}^*} = \text{proj}(M_t | \underline{X}_t).$$

This SDF is unique in the pay-off space.⁴ We let P_t be the arbitrage-free price of an asset with non-negative stochastic payoffs $\{D_t^i\}$ that are measurable w.r.t z^t :

$$P_t^i = E_t \sum_{\tau=t}^{\infty} \frac{\Lambda_{\tau}^*}{\Lambda_t^*} D_{\tau}^i. \quad (4)$$

⁴This requires free portfolio formation and the law of one price (Cochrane 2001).

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:*

$$P_t^L = E_0 \left[\sum_{t=0}^{\infty} \frac{\Lambda_t^*}{\Lambda_0^*} L_t \right], \quad W_t = E_0 \left[\sum_{t=0}^{\infty} \frac{\Lambda_t^*}{\Lambda_0^*} C_t \right].$$

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:

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

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 long-run risk and external habit models we discuss in Section 5 and the heterogeneous agent model in Appendix A. 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. In Section 3.4 below, we

study *good-deal bounds* on the consumption risk premium (see Cochrane and Saa-Requejo (2000) and Alvarez and Jermann (2004)). 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).⁵ 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-, 3-, 4-, and 5-year excess return 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 only 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

⁵Many 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.

income growth fmp.⁶ 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).⁷

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

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

with $\epsilon_t \sim i.i.d. \mathcal{N}(0, I)$ 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, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}'}.$ $\Sigma^{\frac{1}{2}}$ 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 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 bond market variables, the dividend yield and the market return govern the dynamics of stocks. 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. We impose the same predictability structure on the fmp returns. Taken together, Ψ_{21} is a 4×4 matrix with non-zero in the first two columns. In our benchmark model, consumption and labor income growth do not predict future bond and stock market variables; Ψ_{13} and Ψ_{23} are 4×2 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_{31} and Ψ_{32} are 2×4 matrices of non-zero elements and Ψ_{33} is a 2×2 matrix with one zero in the upper-right corner). This structure allows for rich dynamics

⁶We 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. We follow the same procedure for the labor income growth fmp.

⁷Interestingly, 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.

in expected consumption and labor income growth.⁸ In sum, our benchmark Ψ matrix has the following block-diagonal structure:

$$\Psi = \begin{pmatrix} \Psi_{11} & \mathbf{0} & \mathbf{0} \\ \Psi_{21} & \Psi_{22} & \mathbf{0} \\ \Psi_{31} & \Psi_{32} & \Psi_{33} \end{pmatrix}.$$

In section 4, we explore various alternative restrictions on Ψ . These do not materially alter the dynamics of the estimated wealth-consumption ratio.

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. Likewise, the nominal short rate dynamics satisfy $y_t^s(1) = y_0^s(1) + e'_{yn} z_t$, where $y_0^s(1)$ is the unconditional average nominal short rate 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.

We estimate Ψ by OLS, equation-by-equation, and we form each innovation as follows $z_{t+1}(\cdot) - \Psi(\cdot, \cdot)z_t$. We compute their (full rank) covariance matrix Σ .

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 $M_{t+1}^s = \exp(m_{t+1}^s)$ is conditionally log-normal, where lower case letters continue to denote logs:

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

The real pricing kernel is $M_{t+1} = \exp(m_{t+1}) = \exp(m_{t+1}^s + \pi_{t+1})$.⁹ 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 L_t is assumed to be an affine function of the state:

$$L_t = L_0 + L_1 z_t,$$

for an $N \times 1$ vector L_0 and a $N \times N$ matrix L_1 . $L_{1,11}$ contains the bond risk prices, while $L_{1,21}$ and $L_{1,22}$ contain the stock 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-

⁸Several 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.

⁹It also is 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.

dividend ratio in the stock market enables us to set the element on the second row and fifth column of L_1 equal to zero. Likewise, because the last four variables in the VAR cannot affect expected stock 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. (We have shown that the same SDF that prices the returns on traded assets, also prices the returns on non-traded assets, even in economies with household heterogeneity and segmented markets.) This leads to the following structure for L_1 :

$$L_1 = \begin{pmatrix} L_{1,11} & \mathbf{0} & \mathbf{0} \\ L_{1,21} & L_{1,22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix},$$

where $L_{1,11}$, $L_{1,21}$, and $L_{1,22}$ are 4×4 matrices whose entries are not all zero. We impose corresponding zero restrictions on the mean risk premia in the vector L_0 : $L_0 = [L_{0,1}, L_{0,2}, \mathbf{0}]'$, where $L_{0,1}$ and $L_{0,2}$ are 4×1 vectors and the last two elements are zero. We provide further details on the L_0 and L_1 structure below.

In Section 3.4, we relax the spanning assumption. We derive an upper bound on the consumption risk premium by increasing the risk price for the consumption growth innovation, $L_0(9) > 0$, subject to a good-deal upper bound on the maximum Sharpe ratio.

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*

$$wc_t = A_0^c + A_1^c z_t,$$

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:

$$0 = \kappa_0^c + (1 - \kappa_1^c)A_0^c + \mu_c - y_0(1) + \frac{1}{2}(e_c + A_1^c)' \Sigma (e_c + A_1^c) - (e_c + A_1^c)' \Sigma^{\frac{1}{2}} (L_0 - \Sigma^{\frac{1}{2}} e_\pi) \quad (6)$$

$$0 = (e_c + e_\pi + A_1^c)' \Psi - \kappa_1^c A_1^c - e'_{yn} - (e_c + e_\pi + A_1^c)' \Sigma^{\frac{1}{2}} L_1. \quad (7)$$

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 L_0 and L_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:

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

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

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:

$$\begin{aligned} E_t[r_{t+1}^{c,e}] &\equiv E_t[r_{t+1}^c - y_t(1)] + \frac{1}{2} V_t[r_{t+1}^c] = -Cov_t[r_{t+1}^c, m_{t+1}] \\ &= (e_c + A_1^c)' \Sigma^{\frac{1}{2}} \left(L_0 - \Sigma^{\frac{1}{2}'} e_\pi \right) + (e_c + A_1^c)' \Sigma^{\frac{1}{2}} L_1 z_t \end{aligned} \quad (10)$$

where $E_t[r_{t+1}^{c,e}]$ 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; it solves 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; it solves 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:

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

Despite the restrictions we impose on Ψ and L_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 large consumption risk premia. Such a high consumption risk premium would arise when the covariances between consumption growth and the other aggregate shocks are large and/or when the unconditional risk prices in L_0 that

we estimate from the prices of traded assets are sufficiently large. Clearly, a low estimate of the consumption risk premium and hence a high wealth-consumption ratio are not a foregone conclusion. In fact, in our estimation, we choose L_0 large enough to match the equity premium.

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 L_0 and L_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* $L_{0,1}$ and $L_{1,11}$ 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*, the first two elements of $L_{0,2}$ and elements on the first two rows of $L_{1,21}$ and $L_{1,22}$, 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 block*, the last two elements of $L_{0,2}$ and elements on the last two rows of $L_{1,21}$ and $L_{1,22}$, 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, imposing a human wealth share between zero and one, and imposing a good-deal bound. We re-estimate all parameters in L_0 and all 26 parameters in L_1 , starting with the estimates from the third step as starting values.

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 -year

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 L_1 and Ψ , only the risk prices in the bond block of L_1 and L_0 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:

$$\frac{-1}{20} A^{\$}(20) = y_0^{\$}(20), \tag{11}$$

$$\frac{-1}{20} (B^{\$}(20))' = (e_{yn} + e_{spr})'. \tag{12}$$

Equation (11) imposes that the model matches the unconditional expectation of the 5-year nominal yield $y_0^{\$}(20)$. This provides one restriction on L_0 . We choose to let it identify the second element $L_0[2]$. 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 $L_{1,11}$. We choose to identify the second row in $L_{1,11}$. We impose the restrictions (11) and (12) by minimizing the summed square distance between model-implied and actual yields.

Second, the CP risk factor, which is a linear combination of forward rates is the first element in our VAR: $CP_0 + e'_{cp} z_t$. 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 the forward rates. By matching the mean of the factor in model and data, we can identify one additional element of L_0 ; we choose $L_0[4]$. By matching the dynamics, we can identify four more elements in $L_{1,11}$, one in each of the first four columns; we choose to identify the fourth row in $L_{1,11}$. We impose the restriction that the

CP factor is equal in model and data by minimizing their summed squared distance. This second set of moments allows us to replicate the dynamics of bond risk premia in the data.

We now have identified two elements (rows) in L_0 (in $L_{1,11}$).¹⁰ The first and third elements (rows) in L_0 (in $L_{1,11}$) contain only zeros.

2.2 Block 2: Stocks

In the second step, we turn to the estimation of the risk price parameters in $L_{1,21}$ and $L_{1,22}$. 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:

$$\begin{aligned} E_t^{VAR}[r_{t+1}^{m,e}] &= r_0^m - y_0(1) + \frac{1}{2}e'_{rm}\Sigma e_{rm} + ((e_{rm} + e_\pi)' \Psi - e'_{yn}) z_t, \\ E_t^{SDF}[r_{t+1}^{m,e}] &= e'_{rm}\Sigma^{\frac{1}{2}} \left(L_0 - \Sigma^{\frac{1}{2}'} e_\pi \right) + (e_{rm} + e_\pi)' \Sigma^{\frac{1}{2}} L_1 z_t, \end{aligned}$$

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 L_0 ; we choose $L_0[6]$. Matching the risk premium dynamics allows us to identify the second row in $L_{1,21}$ (4 elements) and the second row in $L_{1,22}$ (2 more elements).¹¹ These six elements in $L_{1,22}$ are all needed because expected returns in the VAR depend on the first six state variables. The fifth element of L_0 and the fifth row of L_1 (the first rows of $L_{1,21}$ and $L_{1,22}$) are all zeroes.

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 L_0 and N additional restrictions on L_1 :

$$\begin{aligned} E_t^{VAR}[r_{t+1}^{fmp,e}] &= r_0^{fmp} - y_0(1) + \frac{1}{2}e'_{fmp}\Sigma e_{fmp} + ((e_{fmp} + e_\pi)' \Psi - e'_{yn}) z_t, \\ E_t^{SDF}[r_{t+1}^{fmp,e}] &= e'_{fmp}\Sigma^{\frac{1}{2}} \left(L_0 - \Sigma^{\frac{1}{2}'} e_\pi \right) + (e_{fmp} + e_\pi)' \Sigma^{\frac{1}{2}} L_1 z_t, \end{aligned}$$

¹⁰While the choice of identifying the second and fourth rows is innocuous, it seems natural to associate the prices or risk with the two traded bond yields (short yield and yield spread).

¹¹Again, choosing to identify the sixth element (row) of L_0 (L_1) instead of the fifth row is an innocuous choice. It is more natural to associate the prices of risk with the traded stock return rather than with the non-traded price-dividend ratio.

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 $L_0[7]$ and $L_0[8]$. Matching the risk premium dynamics allows us to identify the third and fourth row in $L_{1,21}$ (8 elements) and the first and second row in $L_{1,22}$ (4 more elements).

2.4 Over-identifying Restrictions

The stock and bond moments described thus far exactly identify the 31 market price of risk parameters that we free up in L_0 (5) and in L_1 (26). For theoretical reasons as well as for reasons of fit, we impose several additional constraints. To avoid over-fitting, we choose not to free up additional market price of risk parameters so that these constraints constitute over-identifying restrictions.

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.¹²

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 (see Lettau and Van Nieuwerburgh (2007), Ang and Liu (2007), and Binsbergen and Koijen (2007) for a discussion of the present-value constraint). 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)):

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

¹²We 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.

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:

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

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: $\log(P_t^d(\tau)) = A^m(\tau) + B^m(\tau)z_t$. 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. This restriction guarantees that stock prices reflect the present-value of future dividend growth. Imposing this constraint not only affects the price of cash flow risk (the sixth row of L_t) but also the real term structure of interest rates (the second and fourth rows of L_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).¹³ 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.

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:

$$E_t \left[r_{t+1}^{l,e} \right] = (e_{\Delta l} + A_1^l)' \Sigma^{\frac{1}{2}} \left(L_0 - \Sigma^{\frac{1}{2}'} e_\pi \right) + (e_{\Delta l} + A_1^l)' \Sigma^{\frac{1}{2}} L_1 z_t.$$

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 ($\mu_l = \mu_c$).¹⁴ We define the labor income share, lis_t , as the ratio of aggregate labor

¹³Appendix 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.

¹⁴We 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

income to aggregate consumption. 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.

Good-Deal Bound Finally, we impose that the conditional volatility of the log stochastic discount factor $m_{t+1}^{\$}$, evaluated at $z = 0$, remains below 1 (Cochrane and Saa-Requejo 2000). This volatility measures the maximum Sharpe ratio that the model allows for. In the final estimation, that conditional volatility of m is 0.69.

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 5 yields, the mean annual pricing errors range from -18 bp to +61 bp, the volatility of the pricing errors range from 0-60 bp, and the RMSE from 0-134 bp.¹⁵ 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 L_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 of 5-years and less, 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.

stays between 0 and 1 in all our estimations.

¹⁵Note 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 is only half as big as on the 20-year yield.

[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. The expected equity return series and the price-dividend series together imply an expected dividend growth rate series. The latter has a correlation of 20% with expected stock returns, a number similar to what Lettau and Ludvigson (2005) estimate.

[Figure 4 about here.]

As in Ang, Bekaert, and Wei (2007), 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 (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. Separately identifying real rate risk and inflation risk based on term structure data alone is challenging.¹⁶ We do not have

¹⁶Many 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.

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 L_0 and L_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 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. 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.¹⁷ 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 was -46% and -61%, respectively.¹⁸ In contrast, real per capita total wealth changed by -12% and +27%, respectively.

[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. 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.]

¹⁷Real per capita wealth is the product of the wealth-consumption ratio and observed real per capita consumption.

¹⁸This includes households' mutual funds holdings.

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 $(1 + \mu_c)^\tau$.¹⁹ 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.]

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 (2007), 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: $E_t[r_{t+1}^y - r_0^y] = 0, \forall t$. 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 L_0 and L_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 (3.4%) are much

¹⁹The underlying security is a real perpetuity with a certain cash flow which grows at a deterministic consumption growth rate μ_c .

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, the length of a typical career. 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 “career” 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 career human wealth number to the per capita value of residential real estate wealth from the Flow of Funds. Career human wealth 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)).²⁰ Finally, it is reassuring that our non-human wealth measure exceeds the

²⁰Lettau and Ludvigson (2001a)’s measure *-cay* falls during the stock market crashes of 1974 and 2000-02. It

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.

3.4 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 structural asset pricing models of Section 5 and of Appendix A. To the best of our knowledge, there are no equilibrium models in which aggregate consumption is not spanned by the returns on traded assets. 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 9th element of L_0 , the risk price of the non-traded consumption growth shock that is orthogonal to all previous 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 L_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.

The first line reports our benchmark case in which the non-traded consumption risk is not priced. The conditional Sharpe ratio on the one-period ahead consumption strip is only .09. When we set $L_0(9)$ to 0.1, this Sharpe ratio doubles. 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. We set the good-deal bound on the maximal Sharpe ratio ($std_t[m_{t+1}]$) at one, following Cochrane and Saa-Requejo (2000). This is a high upper bound, because it is twice the value of the Sharpe ratio on equities in our sample. If we increase $L_0(9)$ all the way to 0.8, the market price of risk exceeds the good-deal bound of one, and the consumption risk premium increases to 3.58% per annum. This is an increase of 139 basis points compared to our benchmark case. The average wealth-consumption ratio drops to 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 a stock remains valid.²¹

Moreover, the implied price of non-traded consumption cash flow risk seems implausible. The conditional Sharpe ratio on the consumption strip is 0.81 at $L_0(9) = 0.8$, eight times higher than the Sharpe ratio on the traded consumption strip. There is no reason to expect aggregate non-traded consumption risk to be priced so differently from traded consumption risk. After all, whether the

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.

²¹In a different model, Alvarez and Jermann (2004) compute a consumption risk premium of 0.2% per annum with a good-deal upper bound for the risk premium of 1.3% per annum.

aggregate risk is traded or not, it is borne by the average investor. Even though we may not be violating the good-deal bound, we are creating a *great deal* for households that take on such non-traded consumption risk. This extreme Sharpe ratio on non-traded aggregate consumption risk also seems implausible because the non-traded consumption innovations are by construction orthogonal to all of the macroeconomic events that affect the average investor, such as recessions or financial crises.²²

Finally, in order to match the equity premium by increasing the price of non-traded consumption risk, we need an increase in the maximum Sharpe ratio to three times the good-deal bound. Clearly, non-traded consumption risk cannot account for the gap between the consumption and the equity risk premium.

[Table 2 about here.]

3.5 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):

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

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 $Cov[wc_t, r_t^H]$ 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

²²Freeing up $L_0(9)$ 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 $L_0(10)$ 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 $L_0(9)$ and $L_0(10)$.

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 very similar 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 tight link between the consumption claim and interest rates.

4 Robustness Analysis

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_t^m , that the fmp return for consumption is only predicted by pd_t^m and its own lag, and that the fmp return for labor income is only predicted by pd_t^m , the lagged fmp return for consumption, and its own lag. This zeroes out the block Ψ_{21} and simplifies the block Ψ_{22} 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 $L_{1,21}$ 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. Lagged $c - l$ is also allowed to predict future consumption growth so that Ψ_{33} has non-zero elements everywhere. We keep the simplified structure for Ψ_{21} and Ψ_{22} from the previous exercise, but we allow Δc and $c - l$ to predict future stock and fmp returns. That is, we free up the last six elements in Ψ_{23} . 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_{23} , this specification requires six non-zero elements in $L_{1,23}$. 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_{13} and identifies four elements in $L_{1,13}$ ($L_1[2, 9 : 10]$ and $L_1[4, 9 : 10]$). 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 L_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.]

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. Interestingly, the LRR and EH models turn out to have dramatically different implications for the wealth-consumption ratio, at least under their benchmark parameterizations. This differences further strengthens the need for measurement. The average wealth-consumption ratio in the LRR model matches the 87 number we estimate in the data. We compare the models’ wc ratios with the one extracted from the data.²³

²³We do not attempt to formally test the two models, because strictly speaking, the LRR and EH models are not nested by our model. Their state displays heteroscedasticity, which translates into market prices of risk L_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. A previous version of this paper contained simulated method of moments estimation of both models.

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 intertemporally. 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):

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

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

$$\sigma_{t+1}^2 = \sigma^2 + \nu_1(\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1}, \quad (16)$$

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_t^2 . The two state variables x_t and σ_t^2 capture time-varying growth rates and time-varying economic uncertainty.

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

$$wc_t = A_0^{c,LRR} + A_1^{c,LRR} z_t^{LRR}. \quad (17)$$

Appendix D.2 proves the proposition, following Bansal and Yaron (2004), and spells out the dependence of $A_0^{c,LRR}$ and $A_1^{c,LRR}$ 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 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

²⁴This 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).

²⁵Since 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.

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 final 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 ($e^{A_0^{c,LRR} - \log(4)}$). Remarkably, this is the exact same value we estimated in the data. 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:

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

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 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

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

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 $\bar{s} = \log(\bar{S})$. 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 state variable of the EH model.

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

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

²⁷See 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.

and the sensitivity function takes the following form

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

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 $(A_0^{c,EH}, A_1^{c,EH}, \bar{S})$.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. This formulation of the SDF suggests that, for the EH model to matter for asset prices, it needs to alter the properties of the wc ratio in the right way.

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% 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 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

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

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

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. However, by overstating the variability of wc , the benchmark CC model overstates the predictability of the total wealth return. 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. In the LRR model, the consumption claim looks more like a bond whereas in the EH model it looks more like a stock. The wealth-consumption ratio should be a useful diagnostic to further improve these and other DAPMs along some of the total wealth and consumption dimensions that they currently do not capture.

6 Conclusion

We develop a new methodology for estimating the wealth-consumption ratio in the data, based on no-arbitrage conditions 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 what one would conclude from studying the equity component of that portfolio. The consumption claim

³⁰Appendix 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.

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 has the undesirable 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 the cash flow process rather than through 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 (2007): “The Term Structure of Real Rates and Expected Inflation,” *Forthcoming Journal of Finance*.
- 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 Pricing 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., R. GOLDSTEIN, AND P. COLLIN-DUFRESNE (2005): “Portfolio Choice over the Life Cycle in the Presence of ‘Trickle Down’ Labor Income,” University of Minnesota and University of California at Berkeley.

- BINSBERGEN, J., AND R. KOIJEN (2007): “Predictive Regressions: A Present-Value Approach,” .
- 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, 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.
- 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.
- 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.
- JAGANNATHAN, R., AND Z. WANG (1996): “The Conditional CAPM and the Cross-Section of Expected Returns,” *Journal of Finance*, 51, 3–54.
- JORGENSEN, 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 S. VAN NIEUWERBURGH (2007): “The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street,” *Forthcoming Review of Financial Studies*.
- MALLOY, C. (2005): “The Geography of Equity Analysis,” *Journal of Finance*, 60, 719–755.
- MARTIN, I. (2007): “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.

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

Table 2: Non-traded Consumption Risk

The first column reports the market price of risk $L_0(9)$ 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 $\sqrt{L_0' L_0}$. The last column shows the conditional Sharpe ratio on a one-period ahead consumption strip: $(e_c' \Sigma^{1/2} L_0) / \sqrt{e_c' \Sigma^{1/2} \Sigma^{1/2} e_c}$.

$L_0(9)$	cons. risk premium	$E[WC]$	$std_t(m_{t+1})$	SR on strips
0	2.19%	87	0.69	0.09
0.05	2.56%	81	0.70	0.14
0.1	2.37%	75	0.70	0.18
0.5	3.05%	48	0.85	0.54
0.8	3.58%	38	1.06	0.81
1.0	3.94%	33	1.21	0.99
1.5	4.86%	25	1.65	1.44
2.0	5.79%	20	2.11	1.89
3.0	7.68%	15	3.08	2.79

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-1 predicts stocks	18.00%	5.55%	0.55%	5.93	9.6%	61.7%	650.7
c-1 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.

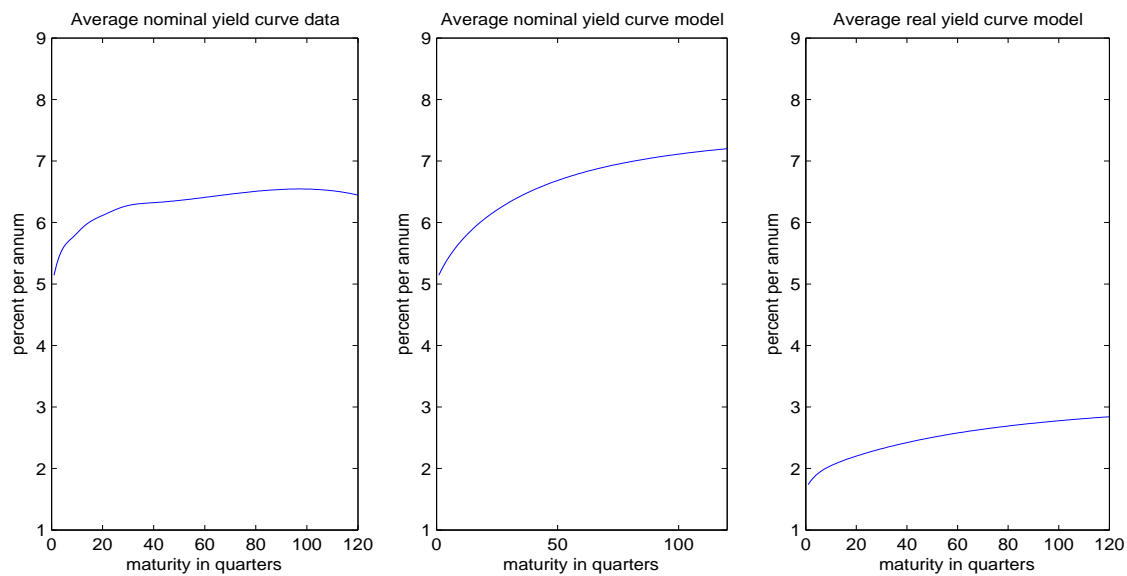


Figure 2: Dynamics of the Nominal Term Structure of Interest Rates

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.

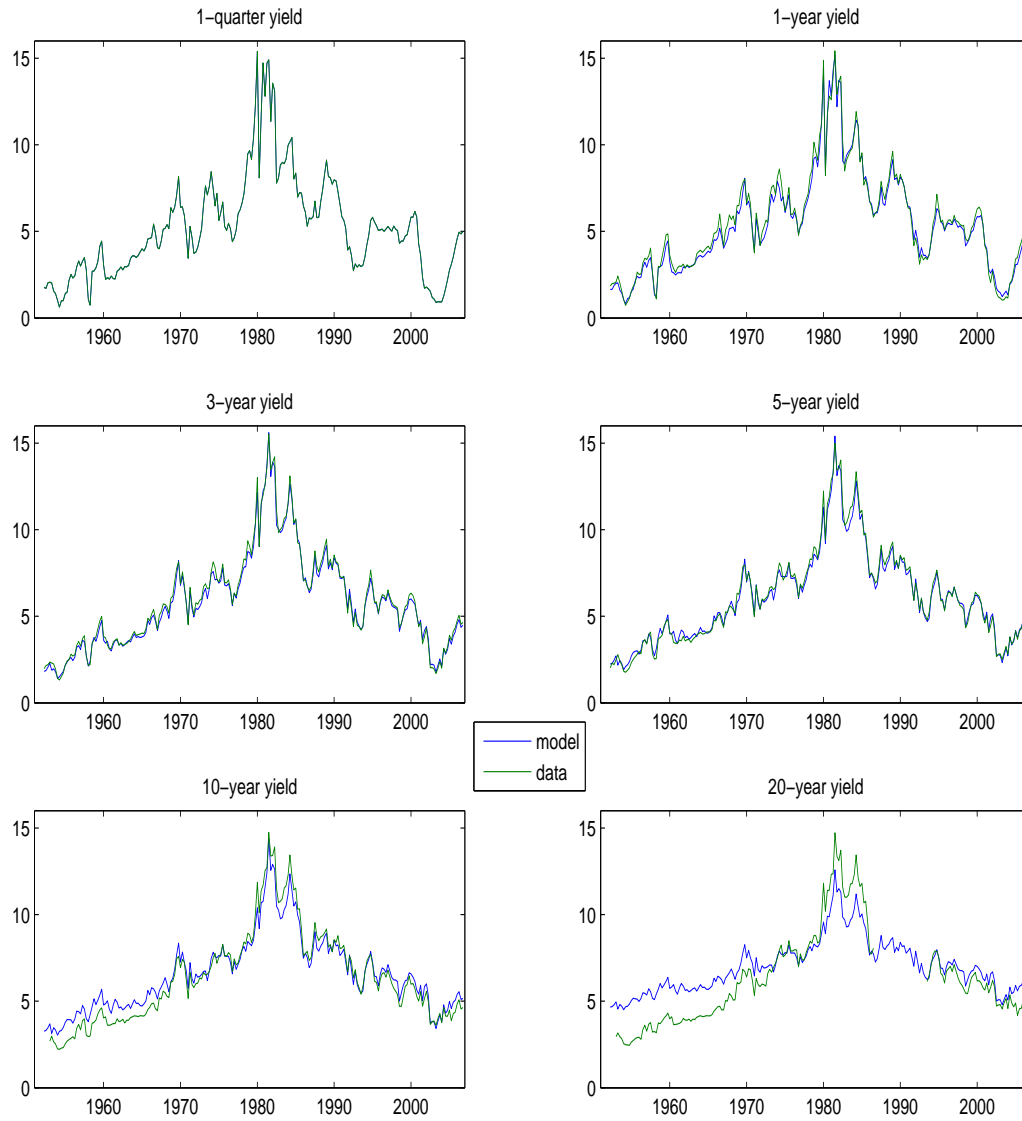


Figure 3: Nominal Bond Risk Premia

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.

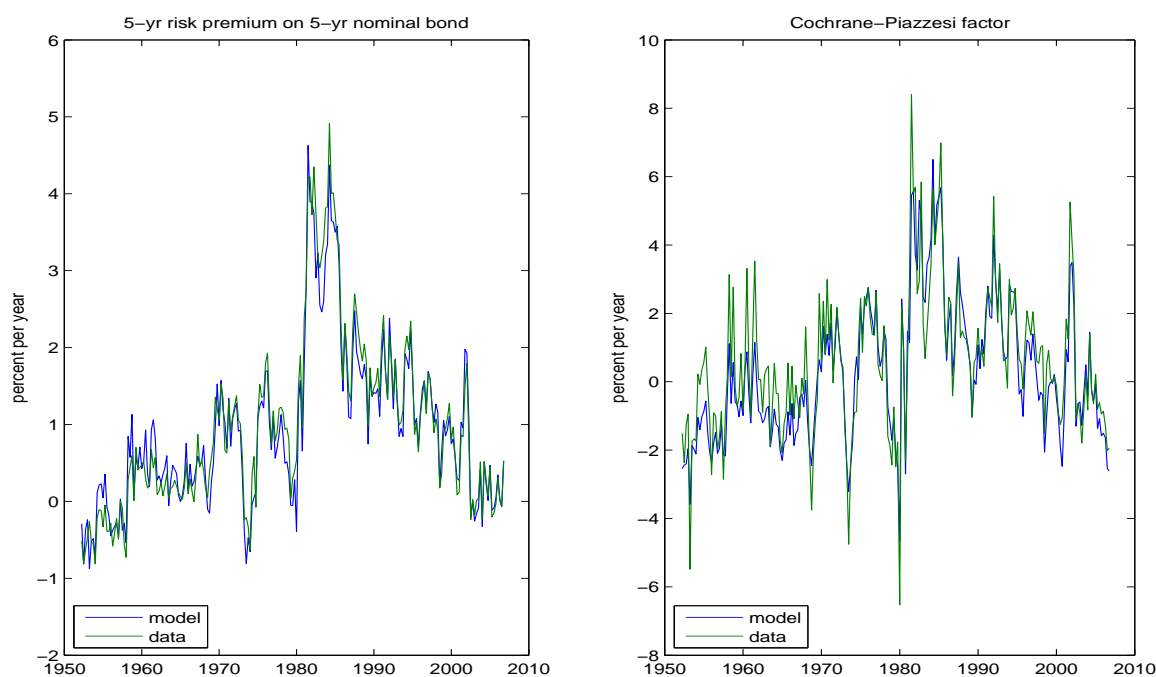


Figure 4: The Stock Market

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

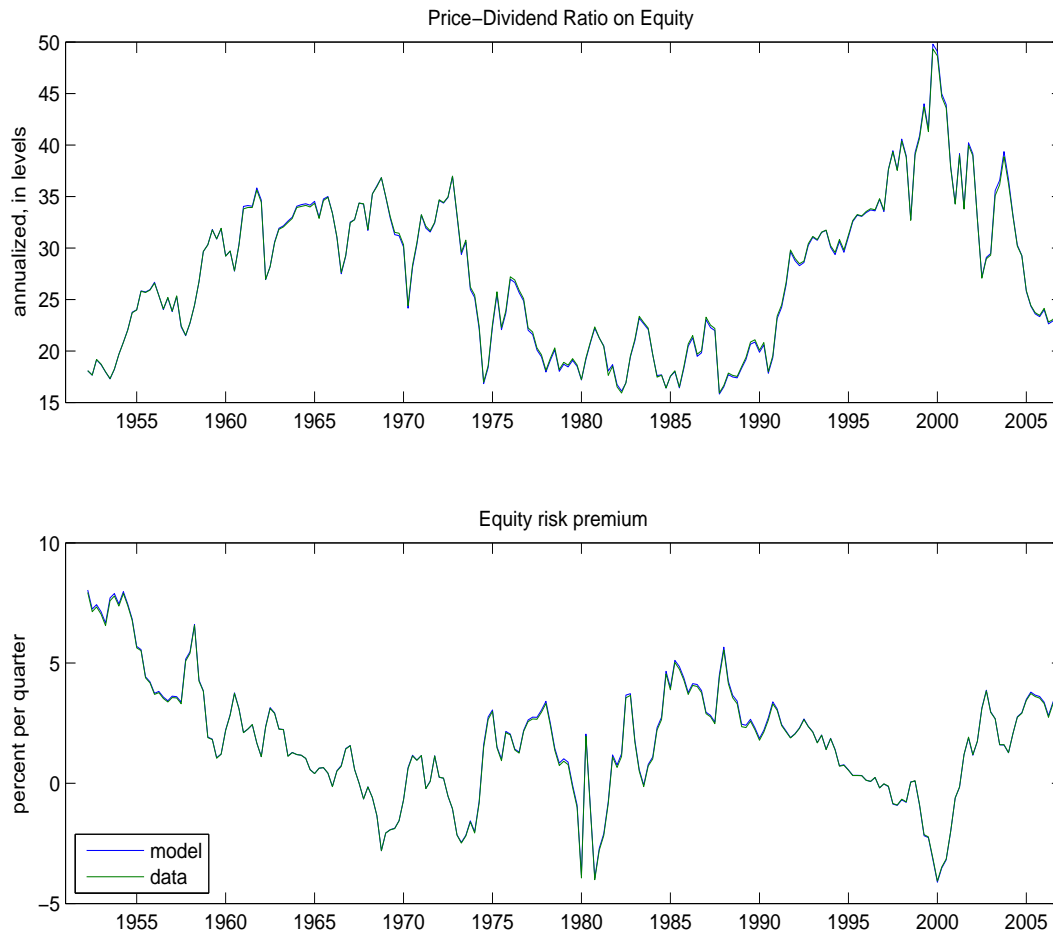


Figure 5: Decomposing the 5-Year Nominal Yield

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.

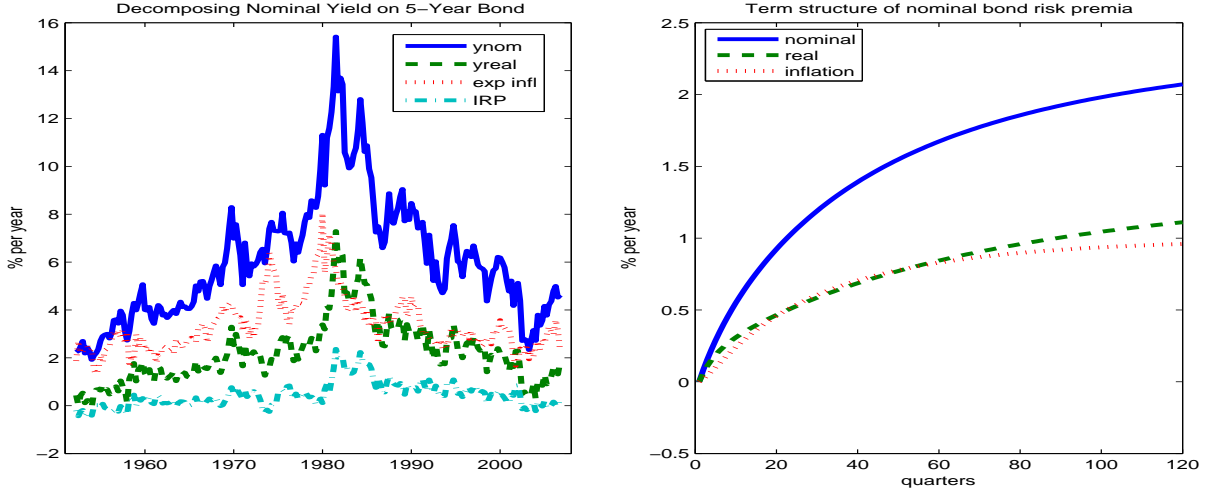


Figure 6: The Log Wealth-Consumption Ratio in the Data

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).

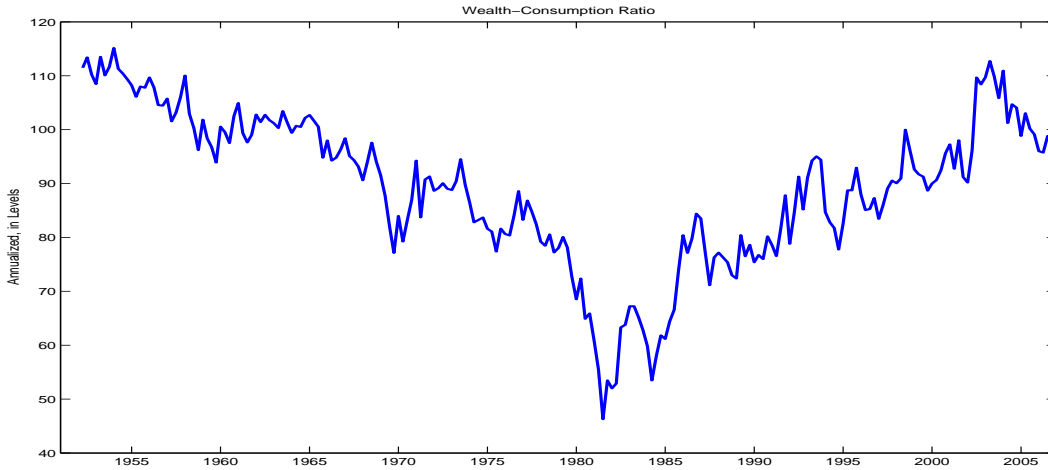


Figure 7: Discount Rates on Consumption and Dividend Claim

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.

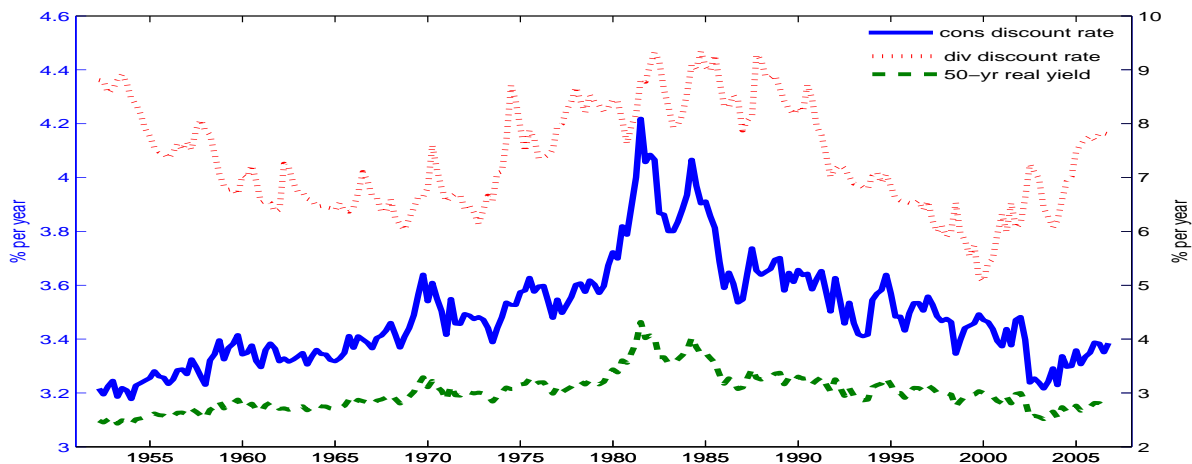


Figure 8: Decomposing the Yield on A Consumption Strip

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.

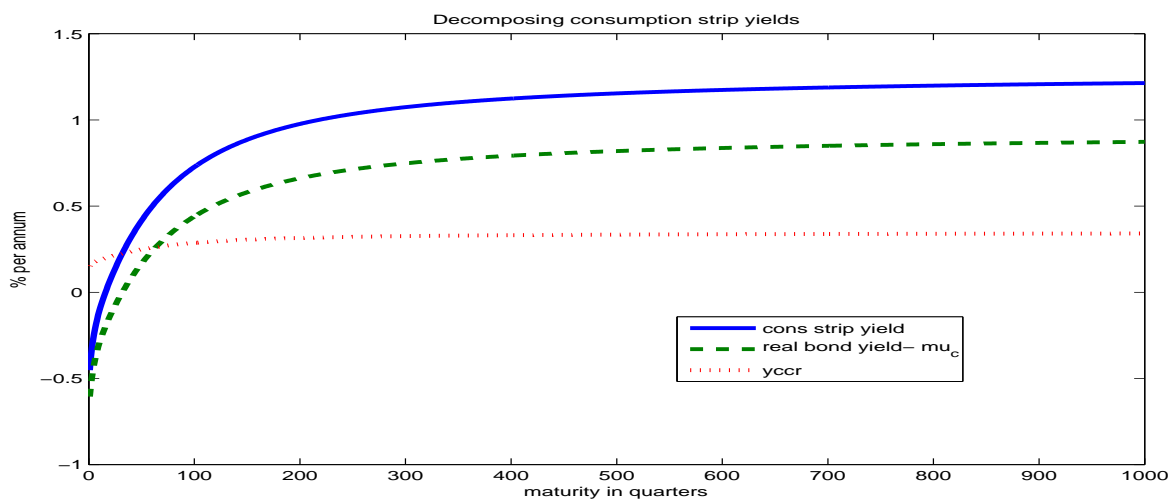


Figure 9: Real Per Capita Wealth Estimates

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 “career human wealth”, the present discounted value of the first 35 years of labor income.

