

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, Robert Goldstein, Lars Peter Hansen, John Heaton, Dana Kiku, Ralph Koijen, Martin Lettau, Francis Longstaff, Sydney Ludvigson, Thomas Sargent, Kenneth Singleton, Stanley Zin, and participants of the NYU macro lunch, seminars at Stanford GSB, NYU finance, BU, the University of Tokyo, LSE, the Bank of England, FGV, MIT Sloan, Purdue, LBS, Baruch, Kellogg, Chicago GSB, and conference participants at the SED in Prague, the CEPR meeting in Gerzensee, the EFA meeting in Ljubljana, the AFA and AEA meetings in New Orleans, the NBER Asset Pricing meeting in Cambridge, and the NYU Five Star Conference 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 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 January 2010
JEL No. E21,G10,G12

ABSTRACT

We set up an exponentially affine stochastic discount factor model for bond yields and stock returns in order to estimate the prices of aggregate risk. We use the estimated risk prices to compute the no-arbitrage price of a claim to aggregate consumption. The price-dividend ratio of this claim is the wealth-consumption ratio. 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, the average US household has more wealth than one might think; most of it is human wealth. A large fraction of the variation in total wealth can be traced back to changes in long-term real interest rates. Contrary to conventional wisdom, we find that events in bond markets, not stock markets, matter most for understanding fluctuations in total wealth.

Hanno Lustig
UCLA Anderson School of Management
110 Westwood Plaza, Suite C413
Los Angeles, CA 90095-1481
and NBER
hlustig@anderson.ucla.edu

Adrien Verdelhan
MIT Sloan School of Management
Department of Finance
50 Memorial Drive, E52-436
Cambridge, MA 02142
and NBER
adrienv@mit.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, stock market wealth is only 2 percent of total household wealth according to our estimates. Real estate, non-corporate businesses, other financial assets, durable consumption goods, and especially human wealth constitute the bulk of total household wealth. The conventional view in asset pricing is that stock market returns are a good proxy for the returns on that remaining 98 percent of the total wealth portfolio (Stambaugh 1982). Much of the empirical work in asset pricing uses the return on the stock market as a stand-in for the total wealth return. Ever since the seminal work of Mehra and Prescott (1985), much of the theoretical work in the equity premium literature models equity as a claim to aggregate consumption. Our results challenge this conventional approach. To understand variation in total wealth, we should look at bond markets, not stock markets.

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. Our measurement does not rely on specifying the preferences of agents, only on the dynamics of aggregate consumption growth. To value this claim, we estimate from stock returns and bond yields the prices of aggregate risk that US households face. We assume that the aggregate consumption growth innovations that are not spanned by traded assets, are not priced.

US households have more wealth than one might think. Our estimates imply that the representative agent 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. The timing of the 1979-81 wealth destruction did not coincide with the stock market crash of 1973-74. Likewise, total wealth was hardly affected by the spectacular decline in the stock market that started in 2000.

We estimate a low consumption risk premium and a high wealth-consumption ratio. The average household's wealth portfolio behaves more like a long-maturity real bond than like equity, for three reasons. First, the total wealth portfolio earns a low risk premium of around 2.2 percent per year, compared to a much higher equity risk premium of 6.9 percent. As a result, the wealth-consumption ratio is much higher, 87 on average, than the price-dividend ratio on equity, 27 on average. Second, the wealth-consumption ratio is less volatile than the price-dividend ratio: its standard deviation is 17 percent versus 27 percent. The return on total wealth has a volatility that is 9.8 percent per year, compared to 16.7 percent for equity returns. Third, the lower variability in the wealth-consumption ratio indicates less variation in expected future total wealth returns. Hence, there is less predictability in total wealth returns than in equity returns. Moreover, most of the variation in future expected total wealth returns is variation in future expected risk-free rates, and not variation in future expected excess returns.

What explains these results is that the consumption risk premium, defined as the risk premium on a claim to stochastically growing aggregate consumption, ends up looking quite similar to the risk premium on a claim to deterministically growing aggregate consumption. The latter claim is a perpetuity with deterministically growing real cash-flows. Its risk premium is the bond risk premium. The risk premia on the risky consumption claim and the real perpetuity are exactly identical if neither innovations to current nor to future consumption growth are priced. These two conditions are approximately satisfied in the data: the consumption cash flow risk accounts for only 15 percent of the consumption risk premium while the bond risk premium accounts for the remaining 85 percent. The resulting spread in risk premia between the two claims is only 0.33 percent per year. These findings are not a foregone conclusion. To allow for priced innovations to future consumption growth, the state vector includes variables that have been shown to predict consumption growth such as the short-term interest rate and the slope of the yield curve. To allow for priced innovations to current consumption growth, we make sure our model prices the returns on a factor mimicking portfolio for consumption growth. Next, we check whether the pricing of consumption innovations that are not spanned by innovations to bond yields or stock returns can overturn our results. Finally, because our results are not driven by the lower volatility of consumption growth, differences between the total wealth portfolio and equity cannot be eliminated by modeling equity as a leveraged claim to aggregate consumption.

Our estimates imply that the real perpetuity with deterministic consumption growth trades at a 37 percent premium to the risky consumption claim. Alvarez and Jermann (2004) show that this premium measures the marginal cost of consumption fluctuations in a large class of representative agent asset pricing models. The representative household would be willing to give up 37 percent of the last unit of annual consumption in exchange for a marginal reduction in aggregate consumption growth risk. This number, which is estimated without taking a stance on preferences, is useful in evaluating the plausibility of structural asset pricing models. After all, it reflects the representative agent's attitudes towards consumption risk which are at the heart of these models. More generally, because the primitives in the leading representative agent asset pricing models are the preferences and the dynamics of aggregate consumption growth, the moments of returns on the consumption claim should be the most informative for testing these models. In contrast, the dividend growth dynamics of stocks can be altered without affecting equilibrium allocations or prices of any traded assets other than stocks; modeling them entails more degrees of freedom.¹ Hence, the fact that several competing models with different primitives match the same moments of equity returns should not come as a surprise. Our approach is to study the *primitive asset pricing moments* of these models and to compare them to the same moments we measure in the data.²

¹Although many model equity as a levered up claim to aggregate consumption, following Abel (1999), they fail to impose restrictions on the long-run behavior of dividends relative to consumption.

²Our goal is not to statistically test these models, since our estimation procedure does not nest them, but simply

The conventional view in modern asset pricing is that equity inherits its risk-return properties from the total wealth portfolio. In this view, equity is risky simply because the total wealth portfolio is. From this perspective, the theoretical challenge is to explain why the price of aggregate consumption growth risk is high and volatile. The first model we study, the external habit model (EH) of Campbell and Cochrane (1999), provides a preference-based answer to this challenge.³ In this model, the consumption and equity claim are similar and the price-dividend ratio on equity and the wealth-consumption ratio are almost identical, even though consumption and dividend growth are not highly correlated.⁴ A version of the EH model that is calibrated to match the equity premium, produces a consumption risk premium of 10.7 percent per annum, while the average wealth-consumption ratio is only 12, 7.25 times smaller than our estimate. Moreover, the wealth-consumption ratio implied by the EH model is almost twice as volatile as the one we have estimated in the data. All of this variation reflects variation in risk premia and not risk-free rates, again inconsistent with our findings. As a result, the EH model dramatically overstates the representative agent's aversion to aggregate consumption fluctuations: the implied marginal cost of consumption risk is 20 times too high.

Our no-arbitrage model prices bonds and stocks rather well, but the estimates lead us to conclude that the total wealth portfolio is more like a portfolio of bonds than stocks. The correlation between realized total wealth returns and stock returns is only 0.12, while the correlation with realized 5-year government bond returns is 0.50. These estimates support the view that equity simply is quite different from the the total wealth portfolio. From our perspective, the new challenge to theorists is to explain why equity is so different from the consumption claim, yet still perceived to be so risky.

The second model we study provides one possible answer to this new challenge. The long-run risk model (LRR) of Bansal and Yaron (2004) makes the dividend claim (equity) much more exposed to shocks to the persistent component of aggregate consumption growth than the consumption claim.⁵ As a result, the LRR model generates a much higher and less volatile wealth-consumption ratio than the price-dividend ratio on equity. The average wealth-consumption ratio

to highlight the key role of the wealth-consumption ratio.

³Early 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 (2008) explores the international finance implications and Chen, Collin-Dufresne, and Goldstein (2008) the implications for credit spreads. Chen and Ludvigson (2007) estimate the habit process for a class of EH models.

⁴See graph on page 220 in Campbell and Cochrane (1999).

⁵Parker and Julliard (2005) and Malloy, Moskowitz, and Vissing-Jorgensen (2009) measure long-run risk based on leads and long-run impulse responses of consumption growth. Bansal, Kiku, and Yaron (2007) estimate the long-run risk model. Piazzesi and Schneider (2006) study its implications for the yield curve, Bansal, Dittmar, and Lundblad (2005) study the implications for the cross-section of equity portfolios, Benzoni, Collin-Dufresne, and Goldstein (2008) for options, and Colacito and Croce (2005) and Bansal and Shaliastovich (2007) for international finance. Chen, Favilukis, and Ludvigson (2008) estimate a model with recursive preferences, while Bansal, Gallant, and Tauchen (2007) estimate both long-run risk and external habit models, and Yu (2007) compares correlations between consumption growth and stock returns across the two models.

in the benchmark LRR model is 87, the same value we estimate in the data, which shows that our numbers are consistent with a standard equilibrium asset pricing model. Most of the variation in the wealth-consumption ratio comes from variability in interest rates, again consistent with the data. However, the LRR mechanism implies more predictability in aggregate consumption growth than is observed in the data.⁶

We turn now to the link to the previous literature. The properties of the average household's total portfolio are crucial for the evaluation of dynamic asset pricing theories. Roll (1977) stresses that the total wealth return is the right pricing factor in the Capital Asset Pricing Model. Similarly, Campbell (1993) shows that current and future total wealth returns substitute for consumption growth as pricing factors in the Intertemporal CAPM. Roll (1977) warns against using return-based factors that proxy for the market return, because if these happen to be ex post mean-variance efficient, the CAPM cannot be rejected, even if it does not hold in the data. However, applied work commonly tests asset pricing models by using the stock market return as a proxy for the total wealth return.

To address this critique, we set out to directly measure the wealth-consumption ratio and the total return on wealth, instead of using proxies. We use a flexible factor model for the stochastic discount factor (henceforth SDF), familiar from the no-arbitrage term structure literature (Duffie and Kan (1996), Dai and Singleton (2000), and Ang and Piazzesi (2003)), and combine it with a vector auto-regression (VAR) for the dynamics of stock returns, bond yields, and consumption and labor income growth, familiar from the methodology of Campbell (1991, 1993, 1996). Like Ang and Piazzesi (2003), we assume that the log SDF is affine in innovations to the state vector, with market prices of risk that are also affine in the same state vector. In a first step we estimate the VAR dynamics of the state. In a second step, we estimate the market prices of aggregate risk. The latter are pinned down by three sets of moments. The first set matches the time-series of nominal bond yields, which are affine in the state, as well as the Cochrane and Piazzesi (2005) bond risk premium. The second set matches the time series of the price-dividend ratio on the aggregate stock market as well as the equity risk premium. We also impose the present value model: the stock price is the expected present-discounted value of future dividends. The third set uses a cross-section of equity returns to form factor-mimicking portfolios for consumption growth and for labor income growth; these are the linear combinations of assets that have the highest correlations with consumption and labor income growth, respectively. We match the time-series of expected excess returns on these two factor-mimicking portfolios. Our SDF model is flexible enough to provide a close fit for the risk premia on bonds and stocks. With the prices of aggregate risk inferred from traded assets, we price the claims to aggregate consumption and aggregate labor income.

⁶Beeler and Campbell (2009) assess the performance of the Long-Run-Risk model and conclude that it implies too much predictability in aggregate consumption growth. However, Bansal and Yaron argue that the model is included in 95 % confidence intervals constructed around the typical stock return moments in the data.

Our methodological approach has two advantages. First, it avoids making arbitrary assumptions about the expected rate of return (discount rate) on human wealth, which is unobserved. Campbell (1993) assumes this discount rate equals the expected rate of return on stocks. Jagannathan and Wang (1996) assume the expected rate of return on human capital is constant, and hence, that the realized rate of return is a linear function of aggregate labor income growth. Lettau and Ludvigson (2001a, 2001b) construct a measure of the log consumption-wealth ratio *cay* as the cointegration between consumption, household financial wealth, and labor income. Its construction implicitly assumes a constant price-dividend ratio on human wealth. Second, our approach avoids using data on housing, durable, and private business wealth from the Flow of Funds.⁷ The variables are often measured at book values and with substantial error. There is no simple way around this measurement problem. Instead, we only use frequently-traded, precisely-measured stock and bond price data and infer the conditional market prices of aggregate risk from them. Using these market prices of risk, we obtain sensible wealth measures. We estimate human wealth to be 90 percent of total wealth. This estimate is consistent with Mayers (1972) who first pointed out that human capital forms a major part of the aggregate capital stock in advanced economies, and with Jorgenson and Fraumeni (1989) who also calculate a 90 percent human wealth share.

Recently, there has been increased interest in understanding the connection between the pricing of bonds and the pricing of equity. Our paper fits into this literature. Our approach is closely related to earlier work by Bekaert, Engstrom, and Grenadier (2005) and Bekaert, Engstrom, and Xing (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 (2009) also match moments in stock and bond markets with an affine model. Finally, Alvarez and Jermann (2004) estimate the consumption risk premium in order to back out the marginal cost of consumption fluctuations from asset prices. Their log SDF is linear in aggregate consumption growth and the total market return. The main difference with our paper is twofold. First, they proxy the total wealth return by a variety of stock return portfolios. Second, their model does not allow for time-varying risk premia. They estimate a smaller consumption risk premium of 0.2 percent, and hence a much higher average wealth-consumption ratio. We show that allowing for time-variation in risk premia and matching conditional moments of not only stock but also bond returns raises the estimated consumption risk premium by 2.0 percent and lowers the wealth-consumption ratio substantially.

We start by measuring the wealth-consumption ratio in the data. Section 1 describes the state variables and their law of motion, while Section 2 shows how we pin down the risk price parameters. Section 3 then describes the estimation results. Section 4 studies the properties

⁷Moskowitz and Vissing-Jorgensen (2002) use Flow of Funds data and Survey of Consumer Finances to value private businesses in the US. Lettau and Ludvigson (2001a, 2001b) also use Flow of Funds data to measure household financial wealth.

of the wealth-consumption ratio in the LRR and EH models. Finally, Section 5 shows that the wealth-consumption ratio estimates are robust to different specifications of the state variables.

1 Measuring the Wealth-Consumption Ratio in the Data

Our objective is to estimate the wealth-consumption ratio and the return on total wealth. Section 1.1 describes our framework. Section 1.2 presents two methodologies to compute the wealth-consumption ratio. Section 1.3 links the wealth-consumption ratio to the cost of aggregate consumption risk.

1.1 Model

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

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

with $\varepsilon_t \sim i.i.d. \mathcal{N}(0, I)$ and Ψ is a $N \times N$ matrix. The vector z is demeaned. The covariance matrix of the innovations is Σ . We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2} \prime}$, which has non-zero elements only on and below the diagonal. 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 the $N \times 1$ vector e_c is the column of a $N \times N$ identity matrix that corresponds to the position of Δc in the state vector. We discuss the state vector in detail below. In addition, the nominal 1-quarter rate is $y_t^{\$}(1) = y_0^{\$(1)} + e'_{yn} z_t$, where $y_0^{\$(1)}$ is the unconditional average nominal short rate and e_{yn} selects the second column of the identity matrix. Likewise, $\pi_t = \pi_0 + e'_{\pi} z_t$ is the (log) inflation rate between $t - 1$ and t with unconditional mean π_0 , etc.

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

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

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

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

to be an affine function of the state:

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

for an $N \times 1$ vector Λ_0 and a $N \times N$ matrix Λ_1 . The matrix $\Lambda_{1,11}$ contains the bond risk prices, $\Lambda_{1,21}$ and $\Lambda_{1,22}$ contain the aggregate stock risk prices, and $\Lambda_{1,31}$ and $\Lambda_{1,32}$ the fmp risk prices.

1.2 The Wealth-Consumption Ratio

We turn now to two approaches to the wealth-consumption ratio. The first definition starts off consumption strips and avoids any approximation. The second definition builds on the Campbell (1991) approximation of log returns.

Consumption Strips A consumption strip of maturity τ pays realized consumption at period τ , and nothing in the other periods. If we impose a no-bubble constraint on the total wealth portfolio, we can compute the wealth-consumption ratio as the sum of the price-dividend ratios on consumption strips of all horizons (Wachter 2005):

$$\frac{W_t}{C_t} = e^{wct} = \sum_{\tau=0}^{\infty} P_t^c(\tau), \quad (2)$$

where $P_t^c(\tau)$ denotes the price of a τ period consumption strip divided by the current consumption. The strip's price-dividend ratio satisfies the following recursion:

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

with $P_t^c(0) = 1$. Appendix B formally proves that the log price-dividend ratios on consumption strips are affine in the state vector.

Not surprisingly, if consumption growth is a random walk with zero risk prices for innovations to current and future consumption growth, then consumption strips are priced like real zero-coupon bonds. The consumption strip price-dividend ratios are approximately equal to the prices of real coupon bonds adjusted for growth μ_c . In this special case, all of the variation in the wealth-consumption can be traced back to the real yield curve.

Consumption strips allow for an exact definition of the wealth-consumption ratio, but they call for the estimation of an infinite sum of bond prices. We turn now to an approximate, but more elegant definition of the wealth-consumption ratio. In our empirical work, we check that both methods deliver similar results.

Wealth Returns In our exponential-Gaussian setting, the log wealth-consumption ratio is an affine function of the state variables. To show this result, we start from the aggregate budget

constraint:

$$W_{t+1} = R_{t+1}^c(W_t - C_t). \quad (3)$$

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

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

We use 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 $A_0^c \equiv E[w_t - c_t]$.⁹

$$r_{t+1}^c \simeq \kappa_0^c + \Delta c_{t+1} + w_{t+1} - \kappa_1^c w_t. \quad (4)$$

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

$$\kappa_1^c = \frac{e^{A_0^c}}{e^{A_0^c} - 1} > 1 \quad \text{and} \quad \kappa_0^c = -\log(e^{A_0^c} - 1) + \frac{e^{A_0^c}}{e^{A_0^c} - 1} A_0^c. \quad (5)$$

Proposition 1. *The log wealth-consumption ratio is then approximately a linear function of the (demeaned) state vector z_t*

$$w_t \simeq A_0^c + A_1^c z_t,$$

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:

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

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

In equation (6), $y_0(1)$ denotes the average real one-period bond yield. This result follows from the Euler equation for the linear approximation of the total wealth return in equation (4) and it is derived in detail in Appendix B. We simply conjecture an affine function for the log wealth-consumption ratio, impose the Euler equation and solve for the coefficients of the affine function. This is an approximation only because it relies on the log-linear approximation of returns in equation (4). Furthermore, this log-linearization is the only approximation in our procedure. Once we have estimated the market prices of risk Λ_0 and Λ_1 (see 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).

⁹Throughout, variables with a subscript zero denote unconditional averages.

This is a system of $N + 1$ non-linear equations in $N + 1$ unknowns; it is non-linear because of equation (5) and can easily be solved numerically.

This solution and the total wealth return definition in (4) 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:

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

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 (see equation 6). This is a key object of interest which measures how risky total wealth is. The second mean-zero term governs the time variation in the consumption risk premium (see equation 7).

Growth Conditions Given the no-bubble constraint, there is an approximate link between the coefficients in the affine expression of the wealth-consumption ratio and the coefficients of the strip price-dividend ratios:

$$\exp(A_0^c) \simeq \sum_{\tau=0}^{\infty} \exp(A^c(\tau)) \quad \text{and} \quad \exp(A_1^c) \simeq \sum_{\tau=0}^{\infty} \exp(B^c(\tau)). \quad (11)$$

A necessary condition for this first sum to converge and hence produce a finite average wealth-consumption ratio is that the consumption strip risk premia are positive and large enough in the limit:

$$(e_c + B^c(\infty))' \Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}'} e_\pi \right) > \mu_c - y_0(1) + \frac{1}{2} (e_c + B^c(\infty))' \Sigma (e_c + B^c(\infty)),$$

where $B^c(\infty)$ satisfies: $B^c(\infty)' = (e_c + e_\pi + B^c(\infty))' \Psi - e'_{yn} - (e_c + e_\pi + B^c(\infty))' \Sigma^{\frac{1}{2}} \Lambda_1$. We refer to this as the growth condition. Because average real consumption growth μ_c exceeds the average real short rate $y_0(1)$ in the data, the right-hand side of the growth condition is positive. If all the risk prices Λ_0 are zero, this condition is obviously violated.

We turn now to a key implication of the wealth-consumption ratio dynamics: the preference-free estimate of the marginal welfare cost of aggregate consumption growth risk that applies to the entire class of representative agent DAPM's. In section 4, we will compare these estimates to the costs implied by two leading models that belong to this class.

1.3 Cost of Aggregate Consumption Risk

The wealth-consumption ratio is directly related to the welfare costs of consumption risk in a large class of models. Within the class of representative agent asset pricing models, it allows us to make inference about the risk preferences of the representative agent, even though the wealth-consumption ratio is derived without committing to a model of preferences. Alvarez and Jermann (2004) define the marginal cost of consumption uncertainty by how much consumption the representative agent would be willing to give up at the margin in order to eliminate some consumption uncertainty.¹⁰ A reduction in exposure to aggregate consumption growth risk can be achieved by selling a claim to stochastically growing aggregate consumption and buying a claim to deterministically growing aggregate consumption. We denote trend consumption by C_t^{tr} . The marginal cost of consumption uncertainty is exactly equal to the ratio of the price of a claim to consumption without cash-flow risk to the price of a claim to consumption with cash-flow risk minus one:

$$\varpi_t = e^{c_t + wc_t - c_t^{tr} - wc_t^{tr}} - 1,$$

where we denote the log price-dividend ratio on the claim to trend consumption, a perpetuity with cash-flows that grow at the average real consumption growth rate μ_c , by wc_t^{tr} . The latter is approximately affine in the state variables: $wc_t^{tr} \simeq A_0^{tr} + A_1^{tr'} z_t$ (see Appendix B for a derivation).

We can approximate the marginal cost of consumption risk as $\varpi_t \simeq wc_t - wc_t^{tr}$ if we evaluate the cost at the consumption trend $c_t = c_t^{tr}$. Hence, the unconditional average of the marginal cost of consumption risk equals $\varpi \simeq A_0^c - A_0^{tr}$, the difference between the mean log wealth-consumption ratio and the mean log price-dividend ratio for the perpetuity. As a benchmark for our results below, it is useful to consider a special case of our model in which $wc_t \approx wc_t^{tr}$, and the marginal cost of business cycles is zero. This case obtains under two conditions. First, if aggregate consumption growth is unpredictable, i.e., $e'_c \Psi = \mathbf{0}$, then innovations to future consumption growth are not priced. Second, if prices of consumption risk are zero, i.e., $e'_c \Sigma^{\frac{1}{2}} \Lambda_1 = 0$ and $e'_c \Sigma^{\frac{1}{2}} \Lambda_0 = 0$, then innovations to current consumption are not priced. If neither current nor future consumption

¹⁰The literature on the costs of consumption fluctuations starts with Lucas (1987) who defines the *total* cost of aggregate consumption risk Ω as the fraction of consumption the consumer is willing to give up in order to get rid of consumption uncertainty: $U((1 + \Omega(\alpha))C^{actual}) = U((1 - \alpha)C^{trend} + \alpha C^{actual})$, where $\alpha=0$. Alvarez and Jermann (2004) define the marginal cost of business cycles as the derivative of this cost evaluated at zero $\Omega'(0)$. While the total cost can only be computed by specifying preferences, the marginal cost can be backed out directly from traded assets prices.

growth innovations are priced, then the risky consumption claim becomes a risk-free consumption claim. In this case, the conditional risk premium on the consumption claim is not zero, but it (approximately) equals the risk premium on the real perpetuity:

$$\begin{aligned} E_t [r_{t+1}^{tr,e}] &\equiv E_t [r_{t+1}^{tr} - y_t(1)] + \frac{1}{2}V_t[r_{t+1}^{tr}] = -Cov_t [r_{t+1}^{tr}, m_{t+1}] \\ &\simeq A_1^{tr'}\Sigma^{\frac{1}{2}} \left(\Lambda_0 - \Sigma^{\frac{1}{2}'}e_\pi \right) + A_1^{tr'}\Sigma^{\frac{1}{2}}\Lambda_1 z_t, \end{aligned} \quad (12)$$

Hence there are two sources of risk premia on consumption above and beyond those on the real perpetuity: (i) positive risk prices for current consumption innovations and (ii) positive risk prices innovations to future consumption growth in the case of consumption growth predictability. In the absence of these, the consumption risk premium equals that on the deterministically growing real perpetuity. In the estimation, we include variables in the state vector that have been shown to forecast future consumption growth, like the short-term interest rate, the slope of the yield curve and the price-dividend ratio, to ensure that our procedure allows for priced innovations to future consumption growth. In addition, we make sure that our model prices factor-mimicking portfolios of consumption growth, to allow for priced current innovations to consumption.

2 Estimating the Market Prices of Risk

In order to recover the dynamics of the wealth-consumption ratio and of the return on wealth, we need to estimate the market prices of risk. We detail in this section our estimation procedure. Section 2.1 lists the additional restrictions we impose on our framework. Section 2.2 describes the estimation technique.

2.1 Restrictions

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 (CP), 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. Aggregate consumption is the sum of nondurable and services consumption, which includes housing services consumption, and durable consumption. In what follows, we use lower-case letters to denote natural logarithms. This state variable is observed at quarterly frequency from 1952.I until 2006.IV (220 observations).¹¹ Appendix A 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).

We use a Cholesky decomposition of the covariance matrix, $\Sigma = \Sigma^{\frac{1}{2}}\Sigma^{\frac{1}{2}'}$, which has non-zero elements only on and below the diagonal. The Cholesky decomposition allows us to interpret the shock to each state variable as the sum of a shock that is orthogonal to the shocks of all preceding state variables and the shocks to all the preceding state variables. Consumption and labor income growth are ordered 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.

Our framework enables us to price both future and current consumption growth shocks. First, the state vector include several variables like interest rates (Harvey (1988)), the price-dividend ratio and the slope of the yield curve (Ang, Piazzesi, and Wei (2006)) that have been shown to forecast future consumption growth. Our first-order VAR explains 19 percent of subsequent consumption growth at quarterly and 37 percent at annual frequencies. In principle, the substantial predictability of future consumption growth allows our model to assign a risk premium to future consumption growth innovations and thus create a wedge between the consumption claim and the real, growing perpetuity.

Second, since the aggregate stock market portfolio has a modest 26% correlation with consumption growth, we use additional information from the cross-section of stocks to learn about the consumption and labor income claims. After all, our goal is to price a claim to aggregate consumption and labor income using as much information as possible from traded assets. We use the 25 size- and value-portfolio returns to form a consumption growth factor mimicking portfolio (fmp) and a labor income growth fmp. Matching factor mimicking portfolio returns allows our model to better price current consumption growth innovations. Moreover, in the estimation, we ensure that our model matches the equity premium. Hence, there is no sense in which a low correlation of consumption growth with returns precludes a high consumption risk premium.

With 8 state variables and time-varying prices of risk, our model has many parameters. On the one hand, it potentially offers a rich environment to describe bond and equity returns, even without latent variables. On the other hand, there is the risk of over-fitting the data. To guard against this risk and to obtain stable estimates, we choose to restrict the dynamics of the model. We do so by building on predictability and asset pricing results established in the equity and bond literature.

¹¹Many 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.

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

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

In section 5, we explore various alternative restrictions on Ψ . These do not materially alter the dynamics of the estimated wealth-consumption ratio. We estimate Ψ by OLS, equation-by-equation, and we form each innovation as follows $z_{t+1}(\cdot) - \Psi(\cdot, :)_t z_t$. We compute their (full rank) covariance matrix Σ .

The zero restrictions on Ψ imply zero restriction on the corresponding elements of the market price of risk dynamics in Λ_1 . For example, the assumption that the stock return and the price-dividend ratio on the stock market do not predict the bond variables CP_t , $y_t^{\$}(1)$, π_t , and $y_t^{\$(20)} - y_t^{\$(1)}$ implies that market prices of the bond market shocks cannot fluctuate with the stock market return or the price-dividend ratio. The entries of Λ_1 in the first four rows and the fifth and sixth column must be zero. Likewise, because the last four variables in the VAR cannot affect expected stock and fmp returns, the prices of stock market risk cannot depend on the last four state

variables. Finally, under our assumption that all sources of aggregate uncertainty are spanned by the innovations to the traded assets (the first eight shocks), 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 relax this assumption in section 5.1. Thus, $\Lambda_{1,41}$, $\Lambda_{1,42}$, $\Lambda_{1,43}$, and $\Lambda_{1,44}$ are zero matrices. This leads to the following structure for Λ_1 :

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

We impose corresponding zero restrictions on the mean risk premia in the vector Λ_0 : $\Lambda_0 = [\Lambda_{0,1}, \Lambda_{0,2}, \Lambda_{0,3}, \mathbf{0}]'$, where $\Lambda_{0,1}$ is 4×1 , and $\Lambda_{0,2}$ and $\Lambda_{0,3}$ are 2×1 vectors.

While all zeroes in Ψ lead to zeroes in Λ_1 in the corresponding entries, the converse is not true. That is, not all entries of the matrices $\Lambda_{1,11}$, $\Lambda_{1,21}$, $\Lambda_{1,22}$, $\Lambda_{1,31}$, and $\Lambda_{1,32}$ must be non-zero even though the corresponding elements of Ψ all are non-zero. Whenever we have a choice of which market price of risk parameters to estimate, we follow a simple rule: we associate non-zero risk prices with traded assets instead of non-traded variables. In particular, we set the rows corresponding to the prices of CP risk, inflation risk, and pd^m risk equal to zero.¹² Our final specification has five non-zero elements in Λ_0 and twenty-six in Λ_1 (two rows of four and three rows of six). This specification is rich enough for the model to match the time-series of the traded asset prices that are part of the state vector. In Section 5.1, 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 in $\Lambda_{0,4} > 0$.

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 B, 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 Λ_t , both the cash flow component and the discount rate component depend on all the stock and the bond components of the state. In the case of cash flows, this follows from the fact that expected consumption growth depends on all lagged stock and bond variables in the state. In the case of discount rates, there is additional dependence through A_1^c , which itself is a function of the first nine state variables. The cash flow component does not

¹²However, just to give an example, the price of stock market risk depends on the price-dividend ratio pd_t^m .

directly depend on the risk prices (other than through κ_1^c) while the discount rate component depends on all risk prices of stocks and bonds through A_1^c . This flexibility implies that our model can accommodate a large consumption risk premium; when the covariances between consumption growth and the other aggregate shocks are large and/or when the unconditional risk prices in Λ_0 are sufficiently large. In fact, in our estimation, we choose Λ_0 large enough to match the equity premium. A low estimate of the consumption risk premium and hence a high wealth-consumption ratio are not a foregone conclusion.

In the benchmark case in which log consumption is a random walk, the cash flow part drops out ($\Delta c_t^H = 0$) and all of the variation in the wealth-consumption ratio is due to variation in returns. If in addition consumption innovations are not priced, then the expected return component equals that of a real perpetuity:

$$r_t^H = [A_1^{tr'}\Psi - \kappa_1^{tr}A_1^{tr'}] (\kappa_1^{tr}I - \Psi)^{-1}z_t.$$

2.2 Estimation

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

We detail the first three steps in Appendix B. They allow us to obtain efficiently the initial vales needed to estimate the full model. In a first step, we estimate the risk prices in the *bond market block* $\Lambda_{0,1}$ and $\Lambda_{1,11}$ by by matching the time series for the short rate, the slope of the yield curve and the CP risk factor. Because of the block diagonal structure, we can estimate these separately. It is key to the actual estimation speed and feasibility that the model delivers bond yields that are affine function of the state variables. In a second step, we estimate the risk prices in the *stock market block* $\Lambda_{0,2}$, $\Lambda_{1,21}$, and $\Lambda_{1,22}$ jointly with the bond risk prices, taking the estimates from the first step as starting values. Here, we impose that the model delivers expected excess stock returns similar to the VAR. In a third step, we estimate the fmp risk prices in the *factor mimicking portfolio block* $\Lambda_{0,3}$, $\Lambda_{1,31}$, and $\Lambda_{1,32}$ taking as given the bond and stock risk prices. Again, we impose that the risk premia on the fmp coincide between the VAR and the SDF model.

The stock and bond moments used in the first three steps exactly identify the 5 elements of Λ_0 and the 26 elements of Λ_1 . In other words, given the structure of Ψ , they are all strictly necessary to match the levels and dynamics of bond yields and stock returns. For theoretical as well as for reasons of fit, we impose several additional constraints. To avoid over-parametrization, we choose not to let these constraints identify additional market price of risk parameters. We re-estimate all 5 parameters in Λ_0 and all 26 parameters in Λ_1 , starting with the estimates from the third step. In each of these steps, we also impose the no-bubble constraint (equation 11), which rules out explosive paths for the wealth-consumption ratio that nevertheless satisfy the Euler equation. 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 B. We now provide more detail on the over-identifying restrictions.

We obtain these restrictions from matching additional nominal yields, imposing the present-value relationship for stocks, and imposing a human wealth share between zero and one, and imposing the growth condition on the consumption claim.

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 restrictions help improve the model’s ability to price distant cash-flows. This is important given that the dynamics of the wealth-consumption ratio will 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. None of these additional term structure constraints, however, are binding at the optimum.¹³

Consumption and Dividend Strips 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.¹⁴ Given a no-bubble-constraint for equities, the sum of the price-dividend ratios on dividend strips of all horizons equals the price-dividend ratio (Wachter 2005):

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

¹³We 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.

¹⁴This constraint is not automatically satisfied from the definition of the stock return: $r_{t+1}^m = \kappa_0^m + \Delta d_{t+1}^m + \kappa_1^m pd_{t+1}^m - pd_t^m$. The VAR implies a model for expected return and the expected log price-dividend ratio dynamics, which implies expected dividend growth dynamics through the definition of a return. These dynamics are different from the ones that would arise if the VAR contained dividend growth and the price-dividend ratio instead. The reason is that the state vector in the first case contains r_t and pd_t^m , while in the second case it contains Δd_t^m and pd_t^m . For the two models to have identical implications for expected returns and expected dividend growth, one would need to include pd_{t-1}^m as an additional state variable. We choose to include returns instead of dividend growth rates because the resulting properties for expected returns and expected dividend growth rates are more desirable. For example, the two series have a positive correlation of 20%, a number similar to what Lettau and Ludvigson (2005) estimate. See Lettau and Van Nieuwerburgh (2007), Ang and Liu (2007), and Binsbergen and Koijen (2009) for an extensive discussion of the present-value constraint.

where $P_t^d(\tau)$ denotes the price of a τ period dividend strip divided by the current dividend. Appendix B.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 (1999) for a similar result. Using (13) and the affine structure, we impose the restriction that the price-dividend ratio in the model equals the one in the data by minimizing their summed squared distance. Imposing this constraint not only affects the price of cash flow risk (the sixth row of Λ_t) but also the real term structure of interest rates (the second and fourth rows of Λ_t). Real yields turn out to play a key role in the valuation of real claims such as the claim to real dividends (equity) or the claim to real consumption (total wealth). 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.

In the estimation, we also impose the no-bubble constraint in equation (11) that the wealth-consumption ratio equals the sum of the consumption strip price-dividend ratios.

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 4 in Appendix B.1.) By the same token, the conditional risk premium on the labor income claim is given by:

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

We use μ_l to denote unconditional labor income growth and $e_{\Delta l}$ selects labor income growth in the VAR. We also impose that aggregate labor income grows at the same rate as aggregate consumption ($\mu_l = \mu_c$).¹⁵ We define the labor income share, lis_t , as the ratio of aggregate labor income to aggregate consumption. It is 0.855 on average in our sample. The human wealth share is the ratio of human wealth to total wealth; it is a function of the labor income share and the price-dividend ratios on human and total wealth:

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

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

¹⁵We 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 5, we estimate a model where we impose cointegration between consumption and labor income by including the log consumption-labor income ratio $c - l$ ratio in place of Δl in the state vector. As explained below, we impose that the human wealth share stays between 0 and 1 in all our estimations.

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. We only report the fit obtained in quarterly data.

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. 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 13bp, and the root mean squared error (RMSE) is 13bp. For the other four yields, the mean annual pricing errors range from -18bp to +61bp, the volatility of the pricing errors range from 10 to 60bp, and the RMSE from 12 to 67bp.¹⁶ While these pricing errors are somewhat higher than the ones produced by term-structure models, our model with only 8 parameters in the term structure block of Λ_1 and especially no latent variables does a good job capturing the level and dynamics of long yields. Furthermore, most of the term structure literature prices yields of maturities up to 5 years, while we also price the 10-year and 20-year yields, because these matter for pricing long-duration assets. On the dynamics, the annual volatility of the nominal yield on the 5-year bond is 1.36% in the data and 1.29% in the model.

The model also does a good job capturing the bond risk premium dynamics. The model produces a nice fit between the Cochrane-Piazzesi factor in model and data. The right panel of Figure 1 shows the CP factor in model and data; it is a measure of the 1-year nominal bond risk premium. The annual mean pricing error is -29bp and standard deviation of the pricing error is 83bp. 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%. The figure suggests our model is able to capture the substantial variation in bond risk premia in the data. This is important because the bond risk premium will turn out to constitute a major part of the consumption risk premium.

[Figure 1 about here.]

The model also manages to capture the dynamics of stock returns quite well. The bottom panel of Figure 2 shows that the model matches the equity risk premium that arises from the VAR

¹⁶Note that the largest errors occur on the 20-year yield, which is unavailable between 1986.IV and 1993.II. The standard deviation and RMSE on the 10-year yield are only half as big as on the 20-year yield.

structure. The average equity risk premium (including Jensen term) is 6.90% per annum in the data, and 7.06% in the model. Its annual volatility is 9.54% in the data and 9.62% the model. The top panel shows the dynamics of the price-dividend ratio on the stock market. The model, where the price-dividend ratio reflects the present discounted value of future dividends, replicates the price-dividend ratio in the data quarter by quarter.

[Figure 2 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 3 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 3 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.¹⁷ Separately identifying real rate risk and inflation risk based on term structure data alone is challenging.¹⁸ We do not have long enough data for real bond yields, but stocks are real assets that contain information about the term structure of real rates. They can help with the identification. For example, higher long real yields in the late 1970s-early 1980s lower the price-dividend ratio on stocks, which indeed was low in the late 1970s-early 1980s (top panel of Figure 2). In terms of average real yields, the dotted line in Figure 4 shows yields ranging from 1.74% per year for 1-quarter real bonds to 2.70% per year for 20-year real bonds.

[Figure 3 about here.]

[Figure 4 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 to save space. The annual risk premium on the consumption growth fmp is 0.79% with a volatility of 1.67 in data and model.

¹⁷Inflation expectations in our VAR model have a correlation of 80% with inflation expectations from the Survey of Professional Forecasters (SPF) over the common sample 1981-2006. The 1-quarter ahead inflation forecast error series for the SPF and the VAR have a correlation of 68%. Realized inflation fell sharply in the first quarter of 1981. Neither the professional forecasters nor the VAR anticipated this decline, leading to a high realized real yield. The VAR expectations caught up more quickly than the SPF expectations, but by the end of 1981, both inflation expectations were identical.

¹⁸Many 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.

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%.

To summarize, table 1 provides a detailed overview of the pricing errors on all these assets. Panel A shows the pricing errors on the equity portfolios, Panels B and C the pricing errors on nominal bonds. Panel A shows that the equity risk premium is 16 basis points per year lower in the data than in the model, while the factor mimicking portfolio returns have the same mean. The volatility and RMSE of the pricing errors on the equity risk premium are about 10 basis points per year. Those on the factor mimicking portfolio returns are 18 and 44 basis points, respectively. Panel B shows the pricing errors on nominal bonds that were used in estimation. The three month rate is matched perfectly since it is in the state vector and carries no risk price. The pricing error on the 5-year bond is only 5 basis points on average, with a standard deviation and RMSE of about 13 basis points. Two- and three-year bonds have pricing error volatilities of 16 and 10 basis points per year. The seven-year bond has a RMSE of 17 basis points, the ten-year bond one of 32 basis points. The largest pricing errors occur on bonds of 20- and 30-year maturity. One mitigating factor is that these bonds have some missing data over our sample period, which makes the comparison of yields in model and data somewhat harder to interpret. Another is that there may be liquidity effects at the long end of the yield curve that are not captured by our model (see Vayanos and Vila (2007)). The pricing errors for nominal bond yields that we obtain are larger than in the standard affine term structure literature, which exclusively focuses on the pricing of 1-year to 5-year bonds using latent factors. We jointly price bonds and stocks, use no latent state variables, and include much longer maturity bonds than what is typically done in the literature.

[Table 1 about here.]

3.2 The Wealth-Consumption Ratio

With the estimates for Λ_0 and Λ_1 in hand, it is straightforward to use Proposition 1 and solve for A_0^c and A_1^c from equations (6)-(7). The third column of Table 2 summarizes the key moments of the log wealth-consumption ratio obtained in quarterly data. The numbers in parentheses are small sample bootstrap standard errors, computed using the procedure described in Appendix B.9. 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 mildly positive (.19).

How risky is total wealth compared to equity? According to our estimation, the consumption risk premium (calculated from equation 10) is 54 basis points per quarter or 2.17% per year. This results in a mean wealth-consumption ratio (A_0^c) of 5.86 in logs, or 87 in annual levels ($\exp\{A_0^c - \log(4)\}$). The consumption risk premium is only one-third as big as the equity risk premium of 6.9%. Correspondingly, the wealth-consumption ratio is much higher than the price-dividend ratio on equity: 87 versus 27. A simple back-of-the-envelope Gordon growth model calculation sheds light on the level of the wealth-consumption ratio. The discount rate on the consumption claim is 3.49% per year (a consumption risk premium of 2.17% plus a risk-free rate of 1.74% minus a Jensen term of 0.42%) and its cash-flow growth rate is 2.34%: $87 = 1/((.0349 - .0234)$. Finally, the volatility of the consumption risk premium is 3.3% per year, one-third of the volatility of the equity risk premium. The standard errors on the moments of the wealth-consumption ratio or total wealth return are sufficiently small so that the corresponding moments of the price-dividend ratio or stock returns are outside the 95% confidence interval of the former. The main conclusion of our measurement exercise is that total wealth is (economically and statistically) significantly less risky than equity.

[Table 2 about here.]

If equity is one natural benchmark for the consumption claim, then the consumption perpetuity is the other natural benchmark. We recall that in the special case of i.i.d. consumption growth and zero risk prices for current and future consumption innovations, the consumption risk premium equals the risk premium on the consumption perpetuity. Table 3 reports the same moments as Table 2 but for the consumption perpetuity. We estimate a risk premium on the perpetuity of 46 basis points per quarter or 1.84% per annum. The difference with the consumption risk premium is positive but small: 33bp per year. Hence, the market assigns a small risk price to current and future consumption innovations. Almost all of the 33bp spread is due to the pricing of innovations to future consumption growth. If consumption is a random walk, the annual consumption risk premium is only 2bp above the perpetuity risk premium. Second, the price-dividend ratio on the perpetuity is less volatile than the wealth-consumption ratio (16% versus 22%). The average log wealth-consumption ratio on the perpetuity $A_0^{c,tr}$ is 6.17, compared to 5.86 for the claim to actual consumption, A_0^c . This implies that the growing real perpetuity trades at a 31% premium ($\varpi \simeq A_0^{c,tr} - A_0^c = 6.17 - 5.86$) on average relative to the consumption claim. On average, the marginal cost of business cycles ϖ is 31% per annum, about twice as high as the benchmark estimate in Alvarez and Jermann (2004), yet much smaller than the costs predicted by the leading representative agent asset pricing models discussed below.

To further understand the source of the small difference between the prices of the consumption claim and the consumption perpetuity, we estimate a simpler *bonds only version* of the model in Section 5.3. This model only matches the bond-market moments by choosing the bond risk prices in the first block of Λ_0 and Λ_1 . All other risk prices are set to zero. This simplified model delivers very similar results to our benchmark. This suggests that the small difference between the consumption risk premium and the consumption perpetuity risk premium is driven by current consumption innovations spanned by bond returns, not stock returns, or future consumption innovations driven by bond variables that predict consumption growth, not by stock returns or the dividend-price ratio.

[Table 3 about here.]

Figure 5 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 3). 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.¹⁹ Similarly, the low-frequency decline of the real yield in the twenty-five years after 1981 corresponds to a gradual rise in the wealth-consumption ratio. One striking way to see that total wealth behaves differently from equity is to study it during periods of large stock market declines. During the periods 1973.III-1974.IV and 2000.I-2002.IV, for example, the change in US households' real per capita stock market wealth, including mutual fund holdings, was -46% and -61%, respectively. In contrast, real per capita total wealth changed by -12% and +27%, respectively. Over the full sample, the total wealth return has a correlation of only 12% with the value-weighted real CRSP stock return, while it has a correlation of 46% with realized one-year holding period returns on the 5-year nominal government bond.²⁰ Likewise, the (1-period) consumption risk premium has a correlation of 22% with the (1-period) equity risk premium, but 66% with the (1-period) nominal bond risk premium.

[Figure 5 about here.]

To show more formally that the consumption claim behaves like a real bond, we compute the discount rate that makes the current wealth-consumption ratio equal to the expected present discounted value of future consumption growth (See Appendix C.2 for details). This is the solid line measured against the left axis of Figure 6. 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

¹⁹Real per capita wealth is the product of the wealth-consumption ratio and observed real per capita consumption.

²⁰A similarly low correlation of 12% is found between total wealth returns and the Flow of Fund's measure of the growth rate in real per capita household net worth, a broad measure of financial wealth. The correlation of the total wealth return with the Flow of Fund's growth rate of real per capita housing wealth is 0.11.

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 large divergence between the perceived riskiness of a claim to consumption and a claim to dividends in securities markets.

[Figure 6 about here.]

Consumption Strips A different way of showing that the consumption claim is bond-like is to study yields on consumption strips. It is useful to decompose the yield on the period- τ strip in two pieces. The first component is the yield on a security that pays a certain cash flow $(1 + \mu_c)^\tau$. The underlying security is a real perpetuity with a certain cash flow which grows at a deterministic consumption growth rate μ_c . The second component is the yield on a security that pays off $C_\tau/C_0 - (1 + \mu_c)^\tau$. It captures pure consumption cash flow risk. Appendix B.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 7 makes clear that the consumption strip yields are mostly comprised of a compensation for time value of money, not consumption cash flow risk.

[Figure 7 about here.]

Potential Limitations Our consumption risk premium estimates could be too low (i) if there is some variable we omitted that predicts consumption growth and carries a positive price of risk, (ii) if we mis-measured the covariance between aggregate consumption growth and the state variables, or (iii) if consumption innovations that are not spanned, are priced. In case (i), we mis-measure the price assigned to innovations to future consumption growth by our model. While we cannot rule out that our VAR misses some predictability of consumption growth at longer horizons, the literature has found modest statistical evidence for consumption growth predictability. Our VAR contains the variables that have been shown to have some predictability: the yield spread, the short-term interest rate, and the price-dividend ratio on the stock market. Our state variables z_t explain 19% of variation in Δc_{t+1} in our benchmark quarterly exercise and even 40% at annual frequency (see Section 5.4). Figure 8 plots the (annualized) one-quarter-ahead expected consumption growth series implied by our VAR. The volatility of annualized expected consumption growth is 0.37% (more than one-third of the volatility of realized consumption growth), while the first-order autocorrelation of expected consumption growth is 0.54 in quarterly data. Moreover, as is clear from the figure, expected consumption growth experiences the largest declines during the severe NBER recessions in the first part of the sample, the 1953.II-1954.II recession, the 1957.III-1958.II recession, the 1973.IV-1975.I recession, the double-dip NBER recession from 1980.I to 1982.IV, and somewhat

smaller declines during the less severe 1960.II-1961.I, 1990.III-1991.I and 2001.I-2001.IV recessions. Hence, the innovations to expected consumption growth are highly cyclical, and could reasonably be expected to be priced in asset markets.

[Figure 8 about here.]

In cases (ii) and (iii), we mis-measure the price assigned to current consumption innovations. To address the second concern, we introduced the factor mimicking portfolios; the consumption (labor income) growth fmp has a 43% (50%) correlation with actual consumption. Our model prices these portfolios very well. Finally, Section 5.1 addresses the pricing of aggregate consumption growth risk that is not spanned.

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% (see last row of Table 2). 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. Jagannathan and Wang (1996) assume that expected returns on human wealth equal the expected labor income growth rate; the resulting price-dividend ratio on human wealth is constant. The construction of *cay* in Lettau and Ludvigson (2001a) makes that same assumption. These models can be thought of as special case of ours, imposing additional restrictions on the market prices of risk Λ_0 and Λ_1 . Our estimation results indicate that expected excess human wealth returns have an annual volatility of 3.7%. This is substantially higher than the volatility of expected labor income growth (0.7%), but much lower than that of the expected excess returns on equity (9.6%). Lastly, average (real) human wealth returns (3.9%) are much lower than (real) equity returns (7.4%), but higher than (real) labor income growth (2.3%) and the (real) short rate (1.7%). In sum, our approach avoids having to make arbitrary assumptions on

unobserved human wealth returns. Our findings do not quite fit any of the assumptions on human wealth returns made in previous work.

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

[Figure 9 about here.]

Finally, we compare non-human wealth, the difference between our estimates for total and for human wealth, with the Flow of Funds series for household net worth. The latter is the sum of equity, bonds, housing wealth, durable wealth, private business wealth, and pension and life insurance wealth minus mortgage and credit card debt. Our non-human wealth series is on average 1.7 times the Flow of Funds series. This ratio varies over time: it is 2.2 at the beginning and at the end of the sample, and it reaches a low of 0.7 in 1973. We chose not to use the Flow of Funds net worth data in our estimation because many of the wealth categories are hard to measure accurately or are valued at book value (e.g., private business wealth). Arguably, only the equity component for publicly traded companies is measured precisely, and this may explain why the dynamics of the household net worth series are to a large extent driven by variation in stock prices (Lettau and Ludvigson (2001a)).²¹ It is reassuring that our non-human wealth measure exceeds the net worth series. After all, our series measures the present discounted value of all future non-labor income.

²¹Lettau and Ludvigson (2001a)’s measure *-cay* falls during the stock market crashes of 1974 and 2000-02. It has a correlation of only 0.16 with our wealth-consumption measure while it has a correlation of 0.37 with the price-dividend ratio on stocks.

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 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 B).

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

Finally, these results do not change if we allow for leverage following Abel (1999)'s approach. We price a claim to C_t^κ , with a leverage ratio of $\kappa = 3$. The average price-dividend ratio on this claim (59) is still much higher than that on equity (27). In addition, the variance decomposition

of the price-dividend ratio on that levered consumption claim is virtually identical to that of the wealth-consumption ratio (not reported in the table). In particular, the covariance with future risk-free rates still accounts for 58% of the variance, while only 36% is accounted for by the covariance with excess returns. In sum, introducing leverage does not close the gap between the consumption claim and the equity claim decomposition.

4 The Wealth-Consumption Ratio in Representative Agent Models

In the class of representative agent asset pricing models, the consumption dynamics and the preferences are the primitive building blocks. These cannot be altered without changing the prices of all assets. Not so for dividend growth dynamics. One can change these without affecting the prices of any other assets. Moreover, these models typically do not impose restrictions on the long-run behavior of dividends. So, the moments of returns on the consumption claims are the *primitive asset pricing moments* of the models. With the answer from the data in hand, we can evaluate their predictions for these primitive asset pricing moments. Moreover, we will compare our preference-free marginal cost of consumption risk estimates to the cost measures implied by these models.

We focus on the long-run risk (LRR) and the external habit (EH) models, each of which has received an enormous amount of attention in the modern asset pricing literature. Just like in the affine model we estimated, the log wealth-consumption ratio is linear in the state variables in each of these two models. We do *not* attempt to formally test the two models, only to point out their implications for the wealth-consumption ratio.²² Interestingly, they have quite different implications for the wealth-consumption ratio. This difference motivates our measurement exercise.

4.1 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 C.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

²²The LRR and EH models are not nested by our model. Their state displays heteroscedasticity, which translates into market prices of risk Λ_t that 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.

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 2. *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 C.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.

We calibrate and simulate the long-run risk model choosing the benchmark parameter values of Bansal and Yaron (2004).²³ Column 1 of Table 2 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 on x is about 5). The persistence of both state variables induces substantial persistence in the wc ratio: its auto-correlation coefficient is 0.91 (0.70) at the 1-quarter (4-quarter) horizon. The change in the wc ratio, which is the second asset pricing factor in the log SDF, has a volatility of 0.90%. Aggregate consumption growth, the first asset pricing factor, has a higher volatility of 1.45%. The correlation between the two asset pricing factors is statistically indistinguishable from zero.

The resulting log total wealth return has a volatility of 1.64% per quarter in the LRR model, again lower than in the data. Low autocorrelation in Δwc and Δc generates low autocorrelation in total wealth returns. The total wealth return has a counter-factually high correlation with consumption growth (+.84) because most of the action in the total wealth return comes from consumption growth. The lower panel reports the consumption risk premium, the expected return on total wealth in excess of the risk-free rate (including a Jensen term). Total wealth is not

²³Since their model is calibrated at monthly frequency but the data are quarterly, we work with a quarterly calibration instead. Appendix C.3 describes the mapping from monthly to quarterly parameters, the actual parameter values, and details on the simulation.

very risky in the LRR model; the quarterly risk premium is 40 basis points, which translates into 1.6% per year. Each asset pricing factor contributes about half of the risk premium. The low consumption risk premium corresponds to a high average wealth-consumption ratio; it is 87 expressed in annual levels ($e^{A_0^{c,LRR} - \log(4)}$). Just as in the data, total wealth is not very risky in the LRR model.

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

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

Appendix C.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.²⁴ 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 would generate the same sign on the covariance with returns as in the data. Alternatively, a positive correlation between innovations to x_t 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 1.²⁵ Therefore, the LRR model generates the wedge between total wealth and

²⁴The heteroscedasticity also affects the risk-free rate component, but without heteroscedasticity there would be no time-variation in risk premia.

²⁵See Appendix C.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 (2005) consider versions of the LRR model with cointegration.

equity we also find in the data.

A final moment of interest is the marginal cost of consumption fluctuations. An issue that arises in both the LRR and EH model is that the price of the consumption perpetuity is either very large or not finite. This happens simply because the benchmark versions of these models do not reproduce bond risk premia we see in the data.²⁶ That is why Table 4 evaluates the cost of consumption risk at shorter horizons, listed in the first column. The last column lists this marginal cost of consumption risk that we estimated in the data at these horizons; the first 100 years account for less than half of the total. The second column of Table 4 is for the LRR model. At each horizon, the LRR model overstates the representative agent’s aversion to aggregate consumption risk by a factor around five.

[Table 4 about here.]

4.2 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.

The results with cointegration are similar and are omitted for brevity.

²⁶Recall that if consumption growth exceeds the short rate, a zero bond risk premium cannot satisfy the no-bubble condition (11). In the benchmark calibration of both models, consumption growth is above the short rate and the bond risk premium is too small. However, work by Bansal and Shaliastovich (2007) on the LRR model and Wachter (2006) on the EH model has shown that these models can be extended and recalibrated to match the term structure of interest rates in the data.

Proposition 3. *The log wealth-consumption ratio is linear in the sole state variable $z_t^{EH} = s_t - \bar{s}$,*

$$wc_t = A_0^{c,EH} + A_1^{c,EH} z_t^{EH}, \quad (19)$$

and the sensitivity function takes the following form

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

Appendix D.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 $(A_0^{c,EH}, A_1^{c,EH}, \bar{S})$.²⁷ 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 D.1 shows this result more formally. Therefore, differences in the properties of the wc ratio between models generates differences in their asset pricing predictions.

We calibrate the EH model choosing the benchmark parameter values of CC.²⁸ 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 2 reports the moments of the wealth-consumption ratio under the benchmark calibration of the EH model. First and foremost, the wc ratio is volatile in the EH model: it has a standard deviation of 29.3%, which is 12.5 times larger than in the LRR model and 12 percentage points higher than in the data. This volatility comes from the high volatility of the surplus consumption ratio (38%). The persistence in the surplus-consumption ratio drives the persistence in the wealth-consumption ratio: its auto-correlation coefficient is 0.93 (0.74) at the 1-quarter (4-quarter) horizon. The change in the wc ratio has a volatility of 9.46%. This is more than 10 times higher than the volatility of the first asset pricing factor, consumption growth, which has a standard deviation of 0.75%.

The high volatility of the change in the wc ratio translates into a highly volatile total wealth return. The log total wealth return has a volatility of 10.26% per quarter in the EH model. As in the LRR model, the total wealth return is strongly positively correlated with consumption growth

²⁷Details are in Appendix D.2. Appendix D.3 discusses an alternative way to pin down \bar{S} .

²⁸Appendix D.4 describes the mapping from monthly to quarterly parameters and reports the parameter values.

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

The properties of total wealth returns are similar to those of equity returns.²⁹ 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 essentially equates the riskiness of total wealth and equity, and as a result, it overstates the representative agent's aversion to consumption risk. The third column of Table 4 shows that the EH model overstates this cost by a factor of at least twenty.

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.³⁰ In the LRR model, as in the data, the consumption claim looks more like a bond, whereas in the EH model it looks more like a stock. Both models overstate the representative agent's aversion to aggregate consumption growth risk, but the EH model is off by much more.

5 Robustness Analysis

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

²⁹Appendix D.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.

³⁰A previous version of this paper contained simulated method of moments estimation of both models, targeting both total wealth and equity return moments.

we consider an estimation at annual frequency.

5.1 Non-Traded Consumption Risk

Sofar we have assumed that all aggregate shocks are spanned by stock and bond prices. This assumption is satisfied in the asset pricing models of Section 4. Even in incomplete markets models, asset prices will reflect changes in the income or wealth distribution (e.g., Constantinides and Duffie (1996)). It is not possible to conclusively bound the wealth-consumption ratio in the absence of spanning, short of writing down a fully specified general equilibrium model. Nevertheless, we want to put some reasonable bounds on the non-traded consumption risk premium in our model. In particular, we relax our assumption that traded assets span all aggregate shocks by freeing up the 9th element of Λ_0 , the risk price of the non-traded consumption growth shock that is orthogonal to the eight traded asset shocks. Table 5 reports the consumption risk premium (Column 2), the average wealth-consumption ratio (Column 3), the maximum conditional Sharpe ratio (Column 4) and the Sharpe ratio on a one-period ahead consumption strip (Column 5) for different values of the price of non-traded consumption risk, governed by the 9th-element of Λ_0 (Column 1). This parameter does not affect the prices of any traded assets, so this exercise does not change any of the model’s implications for observables.³¹

The first line reports our benchmark case in which the non-traded consumption risk is not priced. The consumption risk premium is 2.2% per annum, the maximum Sharpe ratio is 0.7 and the conditional Sharpe ratio on the one-period ahead consumption strip is .09. Increasing $\Lambda_0(9)$ increases the consumption risk premium, lowers the wealth-consumption ratio, and increases the Sharpe ratio on the consumption claim. How far should we increase $\Lambda_0(9)$? A first answer is to bound the maximal Sharpe ratio ($std_t[m_{t+1}]$). Cochrane and Saa-Requejo (2000) and Alvarez and Jermann (2004) choose a “good deal” bound of one, which they argue is high because it is twice the 0.5 Sharpe ratio on equities in the data.³² Since we work with quarterly log returns, the Sharpe ratio on equities is only 0.22, and that same good deal bound of one is more than twice as conservative. This bound is reached for $\Lambda_0(9)$ around 0.8, and implies a consumption risk premium of 3.58% per annum and an average wealth-consumption ratio of 38. Even then, the consumption risk premium is still 4.3% short of the equity premium, so that our conclusion that total wealth

³¹Freeing up $\Lambda_0(9)$ also affects the risk premium and price-dividend ratio on human wealth, in quantitatively similar ways. We also experimented with freeing up the price of risk on the shock to labor income growth that is orthogonal to all previous shocks, including the aggregate consumption growth shock. Increasing this $\Lambda_0(10)$ has no effect on the consumption risk premium and the wealth-consumption ratio. It only affects the risk premium on human wealth. Quantitatively, those effects are similar to those presented in Table 5. The same is true when we simultaneously increase $\Lambda_0(9)$ and $\Lambda_0(10)$.

³²In related work, Bernardo and Ledoit (2000) bound the gain-loss ratio which summarizes the attractiveness of a zero-price portfolio. It is equivalent to a restriction on admissible pricing kernels, precluding the existence of arbitrage and approximate arbitrage opportunities.

has different risk-return characteristics than equity remains valid. In order to match the equity premium by increasing the price of non-traded consumption risk, we would need an increase in the maximum Sharpe ratio to three times the good-deal bound or 14 times the Sharpe ratio on equity.

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

[Table 5 about here.]

5.2 State Dynamics

The results of our estimation exercise are robust to different specifications of the law of motion for the state z . We consider three alternative models. Table 6 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 Ψ_{21} and Ψ_{31} and it simplifies the blocks Ψ_{22} and Ψ_{32} in the companion matrix. Because of the non-zero correlation between the shocks to the term structure and to the stock market variables, the prices of stock market risk inherit an exposure to the term structure variables, so that the elements of $\Lambda_{1,21}$ remain non-zero. The “simple return” specification shows very similar unconditional and conditional moments for the wc ratio. The last column shows a similar fit with the benchmark model; the sum of squared deviations between the moments in the model and in the data is 684 versus 676 in the benchmark. In a second robustness exercise, labeled “c-l predicts stocks”, we replace log labor income growth Δl by the log consumption to labor income ratio $c - l$. This enables us to impose cointegration between the consumption and labor income streams. Just like $E_t[\Delta l_{t+1}]$ before, we assume that $E_t[c_{t+1} - l_{t+1}]$ depends on all VAR elements. In addition, lagged $c - l$ is also allowed to predict future consumption growth so that Ψ_{43} has non-zero elements everywhere. We keep the simplified structure for Ψ_{21} , Ψ_{22} , Ψ_{31} , and Ψ_{32} from the previous exercise, but we allow Δc and $c - l$ to predict future stock and fmp returns. That is, we free up Ψ_{34} and the last two elements in Ψ_{24} .

Consumption growth and to a lesser extent the consumption-labor income ratio have significant predictive power for stock returns and the R^2 of the aggregate return equation increases from 7.6% (benchmark) to 10.6%. This predictability has also been found by Santos and Veronesi (2006) and Lettau and Ludvigson (2001a). Because of the change in Ψ_{24} and Ψ_{34} , this specification requires two (four) additional non-zero elements in $\Lambda_{1,24}$ ($\Lambda_{1,34}$). The third row of Table 6 shows that the wealth-consumption ratio properties are again similar. The mean wealth-consumption ratio is slightly higher and the total wealth return slightly more volatile. The extra flexibility improves the fit. The last exercise, labeled “c-l predicts yield” keeps the structure of the previous robustness exercise, but allows lagged aggregate consumption growth and the lagged consumption-labor ratio to predict the four term structure variables. This frees up Ψ_{14} and identifies four elements in $\Lambda_{1,14}$ ($\Lambda_1[2, 9 : 10]$ and $\Lambda_1[4, 9 : 10]$). The motivation is that a measure of real economic activity, such as consumption growth, is often included as a term structure determinant in the no-arbitrage term structure literature. The wealth-consumption ratio increases a bit further, but not the consumption risk premium. The reason is that the real yield curve is slightly less steep. In conclusion, the various specifications for Ψ and Λ_1 we explored lead to quantitatively similar results. The average consumption risk premium is in a narrow band between 2.16 and 2.24 percent per year; the same is true for the mean wealth-consumption ratio. All calibrations suggest mild predictability of total wealth returns. Whatever predictability there is comes from return predictability, not cash flow predictability. Finally, the future return predictability comes mostly from future risk-free rate predictability, except for the last calibration where risk-free rate predictability is somewhat less pronounced.

5.3 Bond Risk Prices Only

To help us understand what drives the results obtained in the full-fledged model, we re-estimate a simpler version of the model that is nested by the benchmark model: the bonds-only version of the model estimated in step 1 (see section B.7.1). The results are reported in the last row of Table 6. All of the non-bond components of Λ_0 and Λ_1 are set to zero. We estimate the risk prices in the *bond market block* $\Lambda_{0,1}$ and $\Lambda_{1,11}$ by matching only the bond moments: yield time series of maturities 1, 4, 12, 20, 80 quarters (sum of squared deviations), the CP time series (sum of squared deviations), and restrictions on yield curve behavior at very long-horizons. We also impose a good deal bound on the Sharpe ratio of 0.75, and the no-bubble constraint on the wealth-consumption ratio. The results thus obtained are very similar to our benchmark results. The average wealth-consumption ratio is 92.6, compared to 87, and the consumption risk premium is 2.19% per annum, compared to 2.17%. The stock blocks we add in the second and third step of the estimation procedure therefore only have minimal impact on the wealth-consumption estimates. The behavior of the wealth-consumption ratio is almost exclusively determined by the real term

structure, in spite of the fact that the consumption claim is exposed to cash flow risk.³³

[Table 6 about here.]

5.4 Annual Estimation

Annual VAR dynamics may capture lower-frequency correlations between consumption growth and traded asset prices than our benchmark quarterly results. To investigate this possibility, we have re-estimated the wealth-consumption ratio on annual data over the period 1952-2007.³⁴ We find that annual consumption growth has a significantly positive covariance with stock returns (t-stat is 3), which contributes to a better spanning of annual consumption growth risk by the traded assets than in the quarterly model. Nevertheless, the results from the annual estimation, which are listed in the last column of Table 2 are similar to those of the quarterly model. The consumption risk premium is even lower at 1.7% (versus 2.2%) with a volatility of 22% (versus 17%). The dynamics of the wealth-consumption ratio still mirror those of long-term real bond yields. The variance decomposition still attributes all the variance of the wealth-consumption ratio to covariation with future total wealth returns rather than with future consumption growth. There is still some evidence that the wealth-consumption ratio predicts future risk-free rates, albeit weaker than in the quarterly model. Finally, the marginal cost of consumption fluctuations is 168% in the annual model. Our main message that the consumption claim is much less risky than equity remains intact.

6 Conclusion

We develop a new methodology for estimating the wealth-consumption ratio in the data, based on no-arbitrage conditions that are familiar from the term structure literature. Our method combines restrictions on stocks and bonds in a novel way, because we are pricing a claim that a priori has bond-like and stock-like features. However, 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 looks much more like a real bond than like a stock.

³³We obtain almost identical results when we zero out these elements in the Λ_0 and Λ_1 obtained after step 4, instead of re-estimating the model. In this case, the average wealth-consumption ratio is 86.

³⁴The only two differences with the quarterly model is that the short-term rate is now the 1-year Treasury bond yield and that we use the price-dividend ratio on the stock market without repurchase adjustment. The log pd^m ratio has a volatility of 40%. The CP factor is estimated on forward rates constructed from annual yields, and the factor mimicking portfolio returns are also constructed based on annual portfolio returns.

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

A potentially interesting correlation structure could, for example, come out of different firm cohorts. When pricing equity, we typically assume the market portfolio's price-dividend ratio reflects the value of all future dividends rather than the future dividends paid out by the current cohort of firms. However, the valuation of the US stock market in a given month reflects the value of companies that are traded in that month. This cohort of companies will account for an ever decreasing fraction of total dividends paid out in the future.

More generally, by showing that the return on wealth differs substantially for the return on equity, our work raises the bar when it comes to the calibration and estimation of consumption-based dynamic asset pricing models.

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.
- ANG, A., M. PIAZZESI, AND M. WEI (2006): "What does the Yield Curve Tell us about GDP Growth?," *Journal of Econometrics*, 131, 359–403.

- 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 (2007): “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.
- BEELER, J., AND J. Y. CAMPBELL (2009): “The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment,” Working Paper Harvard University.
- 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 (2005): “Risk, Uncertainty and Asset Prices,” Working Paper Columbia University.
- BEKAERT, G., AND S. R. GRENADIER (1999): “Stock and Bond Pricing in an Affine Economy,” NBER Working Paper No. 7346.
- BENZONI, L., P. COLLIN-DUFRESNE, AND R. S. GOLDSTEIN (2008): “Explaining Pre- and Post-1987 Crash Prices of Equity and Options within a Unified General Equilibrium Framework,” Working Paper University of Minnesota.
- BERNARDO, A. E., AND O. LEDOIT (2000): “Gain, Loss and Asset Pricing,” *Journal of Political Economy*, 108 (1), 144–172.
- BINSBERGEN, J., AND R. KOIJEN (2009): “Predictive Regressions: A Present-Value Approach,” *Journal of Finance*, Forthcoming.
- BOUDOUGH, 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. Y. (1991): “A Variance Decomposition for Stock Returns,” *Economic Journal*, 101, 157–179.

- (1993): “Intertemporal Asset Pricing Without Consumption Data,” *American Economic Review*, 83(3), 487–511.
- (1996): “Understanding Risk and Return,” *The Journal of Political Economy*, 104(2), 298–345.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2), 205–251.
- CAMPBELL, J. Y., AND R. J. SHILLER (1988): “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1, 195–227.
- CHEN, L., P. COLLIN-DUFRESNE, AND R. S. GOLDSTEIN (2008): “On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle,” *Review of Financial Studies*, Forthcoming.
- CHEN, X., J. FAVILUKIS, AND S. LUDVIGSON (2008): “An Estimation of Economic Models with Recursive Preferences,” Working Paper New York University.
- CHEN, X., AND S. C. LUDVIGSON (2007): “Land of Addicts? An Empirical Investigation of Habit-Based Asset Pricing Models,” Working Paper New York University.
- COCHRANE, J. H., AND M. PIAZZESI (2005): “Bond Risk Premia,” *American Economic Review*, 95, 138–160.
- COCHRANE, J. H., AND J. SAA-REQUEJO (2000): “Beyond Arbitrage: ‘Good deal’ asset price bounds in incomplete markets,” *Journal of Political Economy*, 108, 79–119.
- COLACITO, R., AND M. CROCE (2005): “Risks for the Long-Run and the Real Exchange Rate,” Working Paper New York University.
- CONSTANTINIDES, G. M. (1990): “Habit-formation: A Resolution of the Equity Premium Puzzle,” *Journal of Political Economy*, 98, 519–543.
- CONSTANTINIDES, G. M., AND D. DUFFIE (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104, 219–240.
- 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.
- 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.
- JORGENSON, D. M., AND B. FRAUMENI (1989): “The Accumulation of Human and Non-Human Capital,” in *The Measurement of Saving, Investment and Wealth*, ed. by R. Lipsey, and H. Tice, pp. 227–282. University of Chicago press, University of Chicago.
- KREPS, D., AND E. L. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46, 185–200.
- LETTAU, M., AND S. C. LUDVIGSON (2001a): “Consumption, Aggregate Wealth and Expected Stock Returns,” *The Journal of Finance*, 56(3), 815–849.
- (2001b): “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying,” *The Journal of Political Economy*, 109(6), 1238–1287.
- LETTAU, M., AND S. C. LUDVIGSON (2005): “Expected Returns and Expected Dividend Growth,” *Journal of Financial Economics*, 76, 583–626.
- LETTAU, M., AND S. VAN NIEUWERBURGH (2007): “Reconciling the Return Predictability Evidence,” Forthcoming Review of Financial Studies.
- LETTAU, M., AND J. WACHTER (2009): “The Term Structures of Equity and Interest Rates,” Working Paper University of California at Berkeley and the Wharton School at the University of Pennsylvania.
- LUCAS, R. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46(6), 1429–54.
- (1987): *Models of Business Cycles*, Yrjo Jahnsson Lecture Series. London: Blackwell.
- 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., T. MOSKOWITZ, AND A. VISSING-JORGENSEN (2009): “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, forthcoming.

- MAYERS, D. (1972): “Nonmarketable Assets in Capital Market Equilibrium,” in *Studies in the Theory of Capital Markets*, ed. by M. C. Jensen, pp. 223–248. Praeger Publishers Inc., New York.
- MEHRA, AND E. PRESCOTT (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15(2), 145–161.
- MENZLY, L., T. SANTOS, AND P. VERONESI (2004): “Understanding Predictability,” *Journal of Political Economy*, 112(1), 1–47.
- MOSKOWITZ, T., AND A. VISSING-JORGENSEN (2002): “The Returns To Entrepreneurial Investment: A Private Equity Premium Puzzle?,” *American Economic Review*, 92, 745–778.
- 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*.
- STAMBAUGH, R. F. (1982): “On The Exclusion of Assets From the Two Parameter Model: A Sensitivity Analysis,” *Journal of Financial Economics*, 10, 237–268.
- VAYANOS, D., AND J.-L. VILA (2007): “A Preferred-Habitat Model of the Term Structure of Interest Rates,” Working Paper London School of Economics.
- VERDELHAN, A. (2008): “A Habit-Based Explanation of the Exchange Rate Risk Premium,” *Journal of Finance*, Forthcoming.
- WACHTER, J. (2005): “Solving Models with External Habit,” *Finance Research Letters*, 2, 210–226.
- (2006): “A Consumption-Based Model of the Term Structure of Interest Rates,” *Journal of Financial Economics*, 79, 365–399.
- YU, J. (2007): “The Long and the Short of Asset Prices: Using Long Run Consumption-Return Correlations to Test Asset Pricing Models,” Working Paper Wharton School of Business.

Table 1: Pricing Errors

This table reports the pricing errors on the asset pricing moments used in the estimation, as well as some over-identifying restrictions. The pricing error time series are computed as the difference between the predicted asset pricing moment by the model and the observed asset pricing moment in the data. The table reports time-series averages (Mean), standard deviations (Stdev), and root-mean squared errors (RMSE). Panel A reports pricing errors on the equity market portfolio, the consumption growth factor-mimicking portfolio (*fmpc*), and the labor income growth factor-mimicking portfolio (*fmpl*). It also reports how well the model matches the price-dividend ratio on the aggregate stock market. Panel B shows nominal bond yield pricing errors for the bond maturities that were used in estimation. Panel C shows bond yield errors for bond maturities that were not used in estimation, as well as the Cochrane-Piazzesi (*CP*) ratio. All moments are annualized and are multiplied by 100, except for the price-dividend ratio which is annualized in levels.

Panel A: Equity Portfolio Returns and PD				
	Equity Mkt.	<i>fmpc</i>	<i>fmpl</i>	pd ratio
Mean	0.1577	-0.0024	0.0047	0.0098
Stdev	0.0786	0.1758	0.4364	0.1601
RMSE	0.1112	0.1754	0.4354	0.1600

Panel B: Nominal Bond Yields Used in Estimation						
	$y^{\$}(1)$	$y^{\$}(4)$	$y^{\$}(12)$	$y^{\$}(20)$	$y^{\$}(40)$	$y^{\$}(80)$
Mean	0.0000	-0.1801	-0.1320	-0.0490	0.2058	0.6119
Stdev	0.0000	0.1624	0.1002	0.1298	0.3010	0.5976
RMSE	0.0000	0.1854	0.1198	0.1318	0.3175	0.6700

Panel C: Nominal Bond Yields Not Used in Estimation and CP					
	$y^{\$}(8)$	$y^{\$}(16)$	$y^{\$}(28)$	$y^{\$}(120)$	<i>CP</i>
Mean	-0.1502	-0.1015	-0.1327	-0.3041	-0.2932
Stdev	0.1379	0.0917	0.1613	0.6132	0.8344
RMSE	0.1567	0.1046	0.1739	0.6289	0.8826

Table 2: 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 two columns report the data at quarterly and annual frequencies respectively. The standard errors are obtained by bootstrap, as described in Appendix B.9.

Moments	LRR Model	EH model	data (quarterly)	data (annual)
$Std[wc]$ (s.e.)	2.35% (.43)	29.33% (12.75)	17.24% (4.30)	22.21% (7.81)
$AC(1)[wc]$ (s.e.)	.91 (.03)	.93 (.03)	.96 (.03)	
$AC(4)[wc]$ (s.e.)	.70 (.10)	.74 (.11)	.85 (.08)	.82 (.21)
$Std[\Delta wc]$ (s.e.)	0.90% (.05)	9.46% (2.17)	4.86% (1.16)	12.17% (3.33)
$Std[\Delta c]$ (s.e.)	1.43% (.08)	.75% (.04)	.44% (.03)	1.07 % (.14)
$Corr[\Delta c, \Delta wc]$ (s.e.)	-.06 (.06)	.90 (.03)	.11 (.06)	0.16 (.16)
$Std[r^c]$ (s.e.)	1.64% (.09)	10.26% (2.21)	4.94% (1.16)	12.44 % (3.42)
$Corr[r^c, \Delta c]$ (s.e.)	.84 (.02)	.91 (.03)	.19 (.07)	0.24 (.15)
$E[E_t[r_t^{c,e}]]$ (s.e.)	0.40% (.01)	2.67% (1.16)	0.54% (.16)	1.72% (.88)
$E[wc]$ (s.e.)	5.85 (.01)	3.86 (.17)	5.86 (.49)	5.02 (.53)
2006 Wealth (in millions) (s.e.)			3.02 (0.27)	4.73 (.52)
hws (s.e.)			0.90 (0.03)	0.90 (.02)

Table 3: Moments of the Consumption Perpetuity

This table displays unconditional moments for the consumption perpetuity, the claim to deterministically growing aggregate consumption. We report its log wealth-consumption ratio wc^{tr} , its first difference Δwc^{tr} , and the log total wealth return $r^{c,tr}$. The last panel reports the time-series average of the conditional consumption risk premium, $E[E_t[r_t^{c,tr,e}]]$, where $r^{c,tr,e}$ denotes the expected log return on total wealth in excess of the risk-free rate and corrected for a Jensen term. We report the estimated moments in the data at quarterly and annual frequencies respectively.

Moments	data (quarterly)	data (annual)
$Std[wc^{tr}]$	16.14	23.70
$AC(1)[wc^{tr}]$	0.95	×
$AC(4)[wc^{tr}]$	0.83	0.83
$Std[\Delta wc^{tr}]$	4.78	12.56
$Corr[\Delta c, \Delta wc^{tr}]$	0.07	0.14
$Std[r^{c,tr}]$	4.79	12.59
$Corr[r^{c,tr}, \Delta c]$	0.07	0.14
$E[E_t[r_t^{c,tr,e}]]$	0.46	1.32
$E[wc^{tr}]$	6.17	6.00
2006 Wealth (in millions)	3,614,122	12,515,109.86

Table 4: Marginal Cost of Consumption Risk

This table displays the marginal cost of consumption fluctuations evaluated at different horizons: $\varpi(H) = \frac{\sum_{\tau=0}^H P_t(\tau)}{\sum_{\tau=0}^H e^{\mu c \tau} P_t^e(\tau)} - 1$.

horizon (years)	LRR model	EH model	data (quarterly)
∞	450,040%		36.59%
10	11.08%	41.31%	1.84%
20	16.60%	82.20%	3.23%
30	22.02%	123.68%	4.73%
50	33.65%	208.45%	7.77%
100	72.02%	431.29%	14.78%

Table 5: Non-traded Consumption Risk

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

$\Lambda_0(9)$	cons. risk premium	$E[WC]$	$std_t(m_{t+1})$	SR on strips
0	2.19%	87	0.69	0.09
0.05	2.26%	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 6: 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 $E[E_t[r_t^{c;\epsilon}]]$ 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	$Std[wc]$	$Std[r^c]$	$E[E_t[r_t^{c;\epsilon}]]$	$E[wc]$	$pred_{CF}$	$pred_{rf}$	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
bond model	18.66%	5.44%	0.55%	5.91	1.4%	64.2%	654.54

Figure 1: 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 2: The Stock Market

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

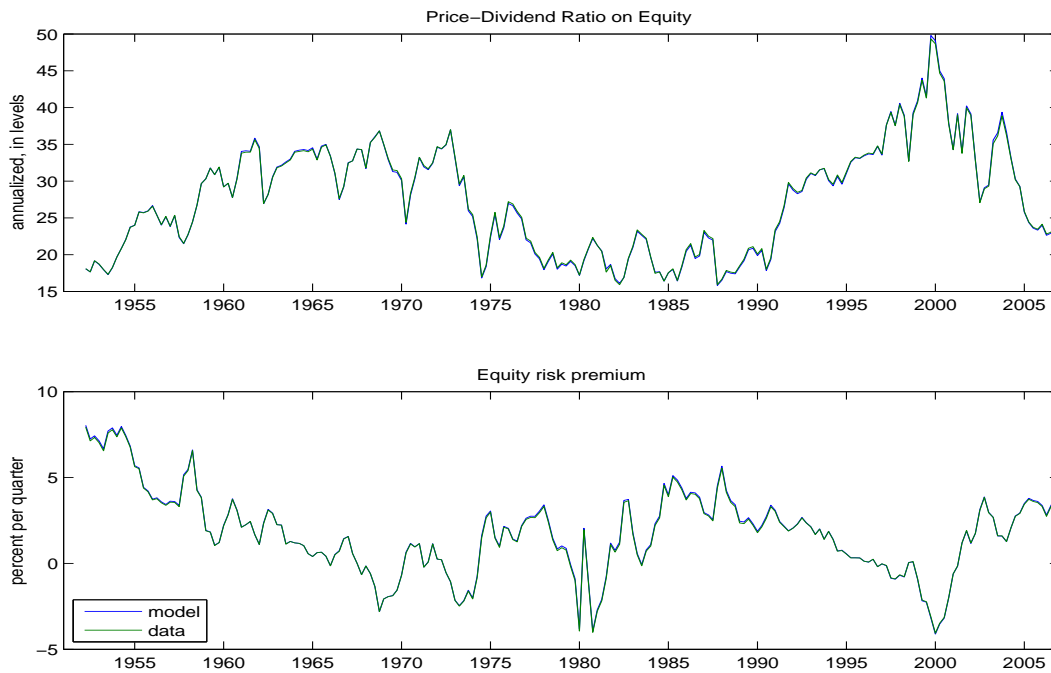


Figure 3: Decomposing the 5-Year Nominal Yield

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 4: 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.

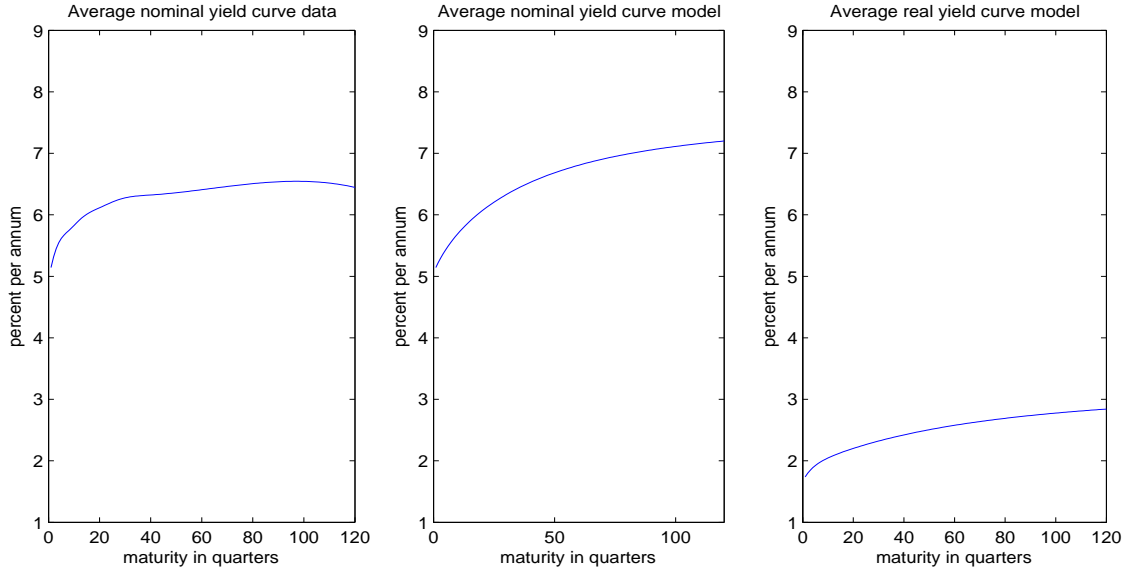


Figure 5: The Log Wealth-Consumption Ratio in the Data

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 6: 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 7: 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 B.5 for a detailed discussion of this decomposition.

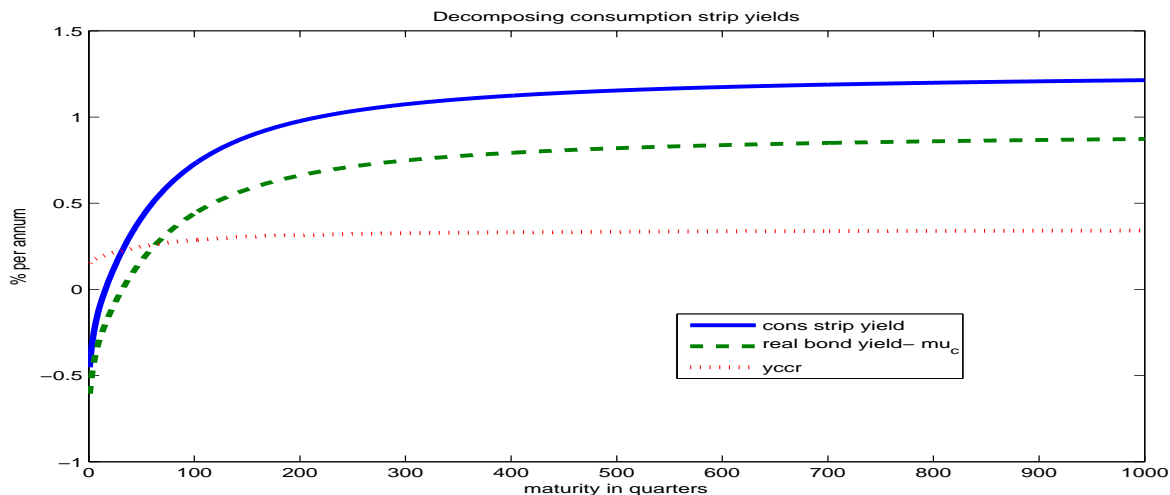


Figure 8: Consumption Growth Predictability

The figure plots (annualized) expected consumption growth at quarterly frequency, as implied by the VAR model: $E_t[\Delta c_{t+1}] = \mu_c + I'_c \Psi z_t$, where z_t is the N -dimensional state vector.

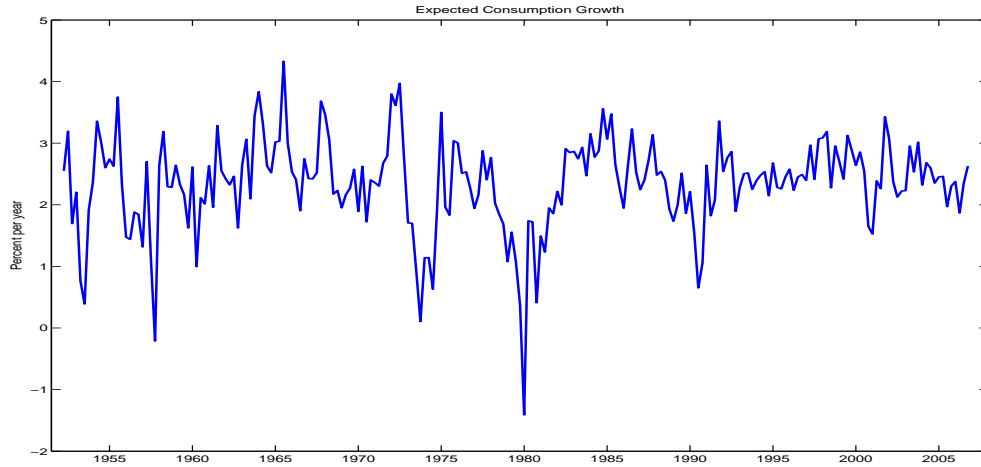


Figure 9: Real Per Capita Wealth Estimates

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

