

NBER WORKING PAPER SERIES

ASSET PRICING WITH COUNTERCYCLICAL HOUSEHOLD CONSUMPTION
RISK

George M. Constantinides
Anisha Ghosh

Working Paper 20110
<http://www.nber.org/papers/w20110>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2014

We thank Lorenzo Garlappi, the late Rick Green, Burton Hollifield, Hanno Lustig, Stijn Van Nieuwerburgh, Thomas Philippon, Bryan Routledge, Ken Singleton (the Editor), Chris Telmer, Sheridan Titman, two anonymous referees, and seminar participants at Carnegie-Mellon University, Concordia University, the EFA 2015 Annual Meeting in Vienna, the ESSFM 2015 conference in Gerzensee, the NBER 2015 summer institute, Georgetown University, New York University, the 2015 Annual Meeting of the SED, the University of British Columbia, the University of California at Los Angeles, the University of Chicago, the University of Miami, the University of Michigan, the University of Southern California, and the University of Texas at Austin for their helpful advice and feedback. Constantinides received financial support from the Center for Research in Security prices, the University of Chicago, as trustee/director of the DFA group of funds, SW7 Holdings, and Cook County Illinois Investment Policy Committee, and as member of the advisory board of FTSE-Russell. Ghosh did not receive financial support from any interested party and does not hold paid or unpaid positions in relevant nonprofit organizations or profit-making entities. The views expressed herein are those of the authors 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by George M. Constantinides and Anisha Ghosh. 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.

Asset Pricing with Countercyclical Household Consumption Risk
George M. Constantinides and Anisha Ghosh
NBER Working Paper No. 20110
May 2014, Revised November 2016
JEL No. D31,D52,E32,E44,G01,G12,J60

ABSTRACT

We show that shocks to household consumption growth are negatively skewed, persistent, countercyclical, and drive asset prices. We construct a parsimonious model where heterogeneous households have recursive preferences. A single state variable drives the conditional cross-sectional moments of household consumption growth. The estimated model fits well the unconditional cross-sectional moments of household consumption growth and the moments of the risk-free rate, equity premium, price-dividend ratio, and aggregate dividend and consumption growth. The model-implied risk-free rate and price-dividend ratio are procyclical while the market return has countercyclical mean and variance. Finally, household consumption risk explains the cross-section of excess returns.

George M. Constantinides
The University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
gmc@ChicagoBooth.edu

Anisha Ghosh
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213
anishagh@andrew.cmu.edu

Considerable empirical evidence documents that households face a substantial amount of uninsurable idiosyncratic labor income risk. The time-variation in idiosyncratic labor income risk plays a central role in understanding several observed phenomena in macroeconomics and finance. Earlier studies focus on the cross-sectional variance of idiosyncratic shocks, arguing that they are countercyclical (e.g., Storesletten, Telmer, and Yaron (2004)) and can, in principle, account for the high historically observed level of the equity premium (e.g., Constantinides and Duffie (1996)). More recently, using a very large data set from the U.S. Social Security Administration, Guvenen, Ozkan, and Song (2014) find that the left skewness of the idiosyncratic shocks is strongly countercyclical. Further, contrary to prior research, they find that the cross-sectional variance is not countercyclical but rather remains mostly flat over the business cycle even after controlling for observable characteristics like age and average past earnings. Krebs (2007) argues that higher job displacement risk in recessions gives rise to the countercyclical left skewness of earnings shocks and that this can generate a substantial cost of business cycles.

In this paper we study the implications of countercyclical left skewness in the cross-sectional distribution of household consumption growth on aggregate asset prices.¹ We construct a parsimonious dynamic equilibrium model with two key ingredients. First, the economy is inhabited by a continuum of heterogeneous households with identical Epstein-Zin (1989) recursive preferences. Second, the heterogeneity among households arises from their labor income processes which are modeled as an exponential function of a Poisson mixture of normal distributions. This specification generates higher-order moments in the cross-sectional distribution of household consumption growth in a tractable fashion. In fact, the parameter driving the Poisson process is the single state variable—hereafter referred to as household

consumption risk—that drives the conditional cross-sectional third central moment of household consumption growth. Aggregate consumption growth is modeled as an i.i.d. process to emphasize that the explanatory power of the model does not derive from such predictability.

We demonstrate that under certain conditions there exists a no-trade equilibrium. To our knowledge, this is the first paper to establish the existence of equilibrium in a heterogeneous-agent, incomplete-markets economy where investors have recursive preferences. For the log-linearized version of the model, we obtain in closed form the equilibrium risk-free rate, expected market return, and price-dividend ratio as functions of the single state variable, household consumption risk.

We estimate the model using the Generalized Method of Moments (GMM) approach, targeting not only the model-implied moments of the market return, the risk-free rate, and the market-wide price-dividend ratio, but also the first three central moments of the cross-sectional distribution of household consumption growth. The estimated model provides a good fit for the time-series averages of the moments of household consumption growth. It also matches well the unconditional mean, volatility, and autocorrelation of the risk-free rate, thereby addressing the risk-free rate puzzle, provides a good fit for the unconditional mean, volatility, and autocorrelation of the market return, thereby addressing the equity premium and excess volatility puzzles, and it matches well the mean, volatility, and autocorrelation of the market price-dividend ratio and aggregate dividend growth, targets that challenge a number of other models. Consistent with empirical evidence, the model implies that the risk-free rate and price-dividend ratio are procyclical while the expected market return, its variance, and the equity premium are countercyclical. The model is also consistent with the salient features of aggregate consumption growth observed in the data: realistic mean and variance, and lack of predictability. Furthermore,

the third central moment of the cross-sectional distribution of household consumption growth explains the cross-section of excess returns of size-, book-to-market equity-, and industry-sorted portfolios as well as the three Fama and French (1993) factors do.

Uninsurable idiosyncratic income shocks is a new paradigm in asset pricing that does not rely on a predictable component in aggregate consumption growth, as in the long-run risks paradigm, time-varying and generally high levels of risk aversion, as in the external habit paradigm, and both large and frequent aggregate macroeconomic disasters, as in the rare disasters paradigm. This paper is the first to empirically establish that observed uninsurable idiosyncratic consumption shocks can explain several observed time-series and cross-sectional patterns in asset market data as well as several aspects of the cross-sectional distribution of household consumption growth.

Figure 1 displays the time series of the skewness in the cross-sectional distribution of quarterly household consumption growth over the period 1982:Q1 to 2009:Q4. The third central moment is mostly negative and countercyclical, with a correlation with NBER recessions between -24.9% and -21.7%. We note that our estimates are noisy because of the measurement error in the survey-based CEX database that we employ in our analysis. Nevertheless, the results suggest that the large countercyclical household income shocks identified in Guvenen, Ozkan, and Song (2014) are transmitted to households as large countercyclical household consumption shocks.

-figure 1 here-

Shocks to household consumption growth are persistent and so are the estimated moments of the cross-sectional distribution of household consumption growth: the autocorrelation of the third central moment is between 12.4% and 18.0%. Similarly, the

autocorrelation of the fifth percentile of the cross-sectional distribution, an alternative measure of skewness, is between 53.4% and 79.2%. The persistent risks play a pivotal role in matching the data, given that the estimated model implies that in the context of recursive preferences, households exhibit a preference for early resolution of uncertainty.

This paper draws on several strands of the literature. First, it builds on the empirical evidence of Attanasio and Davis (1996), Blundell, Pistaferri, and Preston (2008), Cochrane (1991), and Townsend (1994) that consumption insurance is incomplete. Constantinides (1982) highlights the pivotal role of complete consumption insurance, showing that the equilibrium of such an economy with households that have heterogeneous endowments and vonNeumann-Morgenstern preferences is isomorphic to the equilibrium of a homogeneous-household economy. Mankiw (1986) shows that, in a two-period economy with incomplete consumption insurance, the concentration of aggregate shocks among the population is an important determinant of the level of the equity premium. Constantinides and Duffie (1996) further show that, in the absence of complete consumption insurance, given the aggregate income and dividend processes, any given (arbitrage-free) price process can be supported in the equilibrium of a heterogeneous-household economy with judiciously chosen persistent idiosyncratic income shocks. Our paper provides empirical evidence that these shocks are negatively skewed, persistent, and countercyclical, and, more importantly, that they drive asset prices and excess returns.

This paper also draws on Brav, Constantinides, and Geczy (2002) and Cogley (2002), who address the role of incomplete consumption insurance in determining excess returns in the context of economies in which households have power utility. Brav, Constantinides, and Geczy (2002) present empirical evidence that the equity and value premia are consistent over the 1982

to 1996 period, with a stochastic discount factor (SDF) obtained as the average of individual households' marginal rates of substitution with low and economically plausible values of the relative risk aversion (RRA) coefficient. Since these premia are not explained with a SDF obtained as the per capita marginal rate of substitution with low values of the RRA coefficient, the evidence supports the incomplete consumption insurance hypothesis. Cogley (2002) calibrates a model with incomplete consumption insurance that recognizes the variance and skewness of the shocks to households' consumption growth and obtains an annual equity premium of 4.5% to 5.75% with an RRA coefficient of 15. Being framed in terms of economies in which households are endowed with power utility, neither of these papers allows the RRA coefficient and the elasticity of intertemporal substitution (EIS) to be disentangled, a step that appears to be important to address the level and time-series properties of the risk-free rate, price-dividend ratio, and market return.

In contrast to the two papers above, the present investigation disentangles the RRA coefficient and the EIS with recursive preferences. The estimated EIS is low and the model is not subject to the criticism of Dew-Becker and Giglio (2016) and Epstein, Farhi, and Strzalecki (2014) on the extreme implications of models with high EIS regarding the preference for early resolution of uncertainty. In addition, the model addresses the level and time-series properties of the risk-free rate, the price-dividend ratio, and the market return, and it also addresses the cross-section of size-, book-to-market equity-, and industry-sorted portfolio returns.

By introducing recursive preferences, the Euler equations for consumption can no longer be written in terms of household consumption growth alone. It therefore becomes necessary to explicitly model the time-series processes of household consumption and to express the SDF in terms of the consumption-wealth ratio. This complicates model construction and estimation but

has the major side benefit that we no longer need to work with noisy time series of the cross-sectional moments of household consumption growth, but with time-series averages of these moments.

The paper further draws on the literature on macroeconomic crises starting with Rietz (1988) and revisited by Barro (2006) and others as an explanation for the equity premium and related puzzles.² This literature builds on domestic and international evidence that macroeconomic crises are associated with a large and sustained drop in aggregate consumption that increases the marginal rate of substitution of the representative consumer. The basic mechanism of macroeconomic crises in this literature is similar in spirit to that in our paper in that a large decline in the consumption of some or all households increases the marginal rates of substitution of these households. The two classes of models part ways, however, in their quantitative implications. As Constantinides (2008) points out, Barro (2006) calibrates the model by treating the peak-to-trough decline in aggregate consumption during macroeconomic crises (which last on average four years) as if this decline occurs in one year, thus magnifying by a factor of four the size of the observed annual disaster risks. Similar ad hoc magnification of the annual aggregate consumption decline during macroeconomic crises is employed in a number of papers that follow Barro (2006). Using an econometric methodology that allows the probabilities attached to different states of the world to differ from their sample frequencies, and is therefore robust to rare events problem in the data, Julliard and Ghosh (2012) reject the rare events explanation for the equity premium puzzle. They show that to explain the equity premium puzzle with power utility preferences of the representative agent and plausible RRA after the multi-year nature of disasters is taken into account, one should be willing to believe that economic disasters occur every 6.6 years, on average. Moreover, Backus, Chernov, and Martin (2011) demonstrate

that options imply smaller probabilities of extreme outcomes than the probabilities estimated from international macroeconomic data.

In contrast to these models, our model relies on shocks to *household* consumption growth, with a frequency and annual size consistent with empirical observation. These shocks support the observed time-series properties of the risk-free rate, market return, and market price-dividend ratio. Furthermore, the shocks to household consumption “average out” across households and do not imply unrealistically large annual shocks to aggregate consumption growth.

Finally, this paper relates to the literature on the cross-section of excess returns. We show that the third central moment of the cross-sectional household consumption growth distribution explains the cross-section of excess returns on the size-, book-to-market equity-, and industry-sorted portfolios. The results from our one-factor model are comparable to those of the three-factor Fama-French (1993) model.

The paper is organized as follows. The model and its implications for consumption growth and prices are presented in Section I. We discuss the data in Section II. In Section III, we present summary statistics of the cross-sectional distribution of household consumption growth. The empirical methodology and main results are presented in Section IV. We further interpret the results in Section V. In Section VI, we discuss implications of household consumption risk for the cross-section of excess returns. We conclude in Section VII. Derivations are provided in appendices.

I. The Model

We consider an exchange economy with a single nondurable consumption good serving as the numeraire. There are an arbitrary number of traded securities (for example, equities, corporate bonds, default-free bonds, and derivatives) in positive or zero net supply. Conspicuously absent are markets for trading the households' wealth portfolios. A household's wealth portfolio is defined as a portfolio with dividend flow equal to the household's consumption flow. It is in this sense that the market is incomplete, thereby preventing households from insuring their idiosyncratic income shocks. The sum total of traded securities in positive net supply is referred to as the "market." The market pays net dividend D_t at time t , has ex-dividend price P_t , and has normalized supply of one unit. We assume that households are endowed with an equal number of market shares at time 0 but can trade in these shares and all other securities (except the wealth portfolios) thereafter.

Aggregate consumption is denoted by C_t , log consumption by $c_t \equiv \log(C_t)$, and consumption growth by $\Delta c_{t+1} \equiv c_{t+1} - c_t$. We assume that aggregate consumption growth is i.i.d. normal: $\Delta c_{t+1} = \mu + \sigma_a \varepsilon_{t+1}$, $\varepsilon_t \sim \mathcal{N}(0,1)$. By construction, aggregate consumption growth has zero autocorrelation, is unpredictable, and is uncorrelated with business cycles. We also consider the case in which expected growth in aggregate consumption is a function of the state variable that is correlated with the business cycle and obtain similar results. We choose to present the case in which expected growth in aggregate consumption is uncorrelated with the business cycle in order to highlight the role of variation in household consumption risk along the business cycle. Aggregate labor income is defined as $I_t = C_t - D_t$.

There are an infinite number of distinct households and their number is normalized to one. Household i is endowed with labor income $I_{i,t} = \delta_{i,t}C_t - D_t$ at date t , where

$$\delta_{i,t} = \exp \left[\sum_{s=1}^t \left\{ \left(j_{i,s}^{1/2} \sigma \eta_{i,s} - j_{i,s} \sigma^2 / 2 \right) + \left(\hat{j}_{i,s}^{1/2} \hat{\sigma} \hat{\eta}_{i,s} - \hat{j}_{i,s} \hat{\sigma}^2 / 2 \right) \right\} \right]. \quad (1)$$

The exponent consists of two terms inside the summation. The first term captures shocks to household income that are related to the business cycle, for example, job loss by the prime wage-earner in the household. The business cycle is tracked by the single state variable in the economy, $\omega_t > 0$, which follows a Markov process to be specified below. The state variable drives the household income shocks through the random variable $j_{i,s}$, which has a Poisson distribution with $\text{prob}(j_{i,s} = n) = e^{-\omega_s} \omega_s^n / n!$, $n = 0, 1, \dots, \infty$, $E(j_{i,s}) = \omega_s$, and is independent of all primitive random variables in the economy. The random variable $\eta_{i,s} \sim N(0, 1)$ is i.i.d. and independent of all primitive random variables in the economy. Thus, the first term is the sum of variables that are normal, conditional on the realization of $j_{i,s}$. The volatility of the conditional normal variable is $j_{i,s}^{1/2} \sigma$, which is driven by the variable $j_{i,s}$ and has a distribution driven by the state variable.³ The second term of the exponent captures shocks to household income that are unrelated to the business cycle, for example, the death of the prime wage-earner in the household. It is defined in a similar manner as the first term with the major difference being that $\hat{\omega}$ is a parameter instead of a state variable.

This particular specification of household income captures several key features of household income and consumption. First, since the income of the i^{th} household at date t is

determined by the sum of all past idiosyncratic shocks, household income shocks are permanent, which is broadly consistent with empirical evidence that household income shocks are persistent (e.g., Storesletten, Telmer, and Yaron (2004)). Second, the joint assumptions that the number of households is infinite and the income shocks are symmetric across households allow us to apply the law of large numbers and show that the identity $I_t = C_t - D_t$ is respected.⁴ Third, this particular specification of household income, combined with the symmetric and homogeneous household preferences to be defined below, is shown to imply that households choose not to trade in which case household consumption is given simply by $C_{it} = I_{it} + D_t = \delta_{it} C_t$. Finally, the cross-sectional distribution of relative household consumption growth, $\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right)$, has a negative third central moment. Its moments depend on the parameters of the distribution of $j_{i,s}$, which are driven in turn by the state variable ω_s . Hereafter, we refer to the state variable as “household risk.”

We assume that households have identical recursive preferences,

$$U_{i,t} = \left\{ (1-\delta)(C_{i,t})^{1-1/\psi} + \delta \left(E_t \left[(U_{i,t+1})^{1-\gamma} \right] \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{1/(1-1/\psi)}, \quad (2)$$

where δ is the subjective discount factor, γ is the RRA coefficient, ψ is the EIS, and

$\theta \equiv \frac{1-\gamma}{1-1/\psi}$.⁵ As shown in Epstein and Zin (1989), the SDF of household i is

$$SDF_{i,t+1} = \exp\left(\theta \log \delta - \frac{\theta}{\psi} \Delta c_{i,t+1} + (\theta - 1) r_{i,c,t+1}\right), \quad (3)$$

where $\Delta c_{i,t+1} \equiv \log(C_{i,t+1}) - \log(C_{i,t})$ and $r_{i,c,t+1}$ is the log return on the i^{th} household's private valuation of its wealth portfolio. As we discuss in Section V.C, recursive preferences appear to be necessary to explain the equity premium and risk-free rate puzzles.

We conjecture and verify that autarchy is an equilibrium. Autarchy implies that the consumption of household i at date t is $C_{i,t} = I_{i,t} + D_t = \delta_{i,t} C_t$ and household consumption growth $C_{i,t+1} / C_{i,t} = \delta_{i,t+1} C_{t+1} / \delta_{i,t} C_t$ is independent of the household's consumption level.⁶ This feature, combined with the property that the household's utility is homogeneous in the household's consumption level, implies that the return on the household's private valuation of its wealth portfolio is independent of the household's consumption level. The SDF of household i is therefore independent of the household's consumption level—it is specific to household i only through the term $\delta_{i,t+1} / \delta_{i,t}$. In pricing any security, other than households' wealth portfolios, the term $\delta_{i,t+1} / \delta_{i,t}$ is integrated out of the pricing equation and the private valuation of any security is common across households. This result verifies the conjecture that autarchy is an equilibrium. We formalize this argument in Appendix B.

In deriving the result that autarchy is an equilibrium and the equilibrium consumption of household i at date t is $C_{i,t} = \delta_{i,t} C_t$, we rely on the assumption that the market is incomplete and hence prevents households from insuring any component of their idiosyncratic income shocks. A reader who finds implausible the assumption that households cannot insure any component of their idiosyncratic income shocks and the resulting implication that autarchy is an equilibrium

may simply interpret $C_{i,t} = \delta_{i,t} C_t$ as the post-trade consumption of household i . Our empirical methodology is consistent with either of these two interpretations of the relation $C_{i,t} = \delta_{i,t} C_t$ because we use household consumption data and not household income data. The degree of market incompleteness and the relation between household income and household consumption are outside the scope of the present investigation.

The logarithm of cross-sectional relative household consumption growth is

$$\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right) = \delta_{i,t+1} - \delta_{i,t} = j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2.$$

Its conditional central moments, derived in Appendix C, are as follows:

$$\mu_1\left(\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right) \middle| \omega_{t+1}\right) = -\sigma^2 \omega_{t+1} / 2 - \hat{\sigma}^2 \hat{\omega} / 2, \quad (4)$$

$$\mu_2\left(\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right) \middle| \omega_{t+1}\right) = (\sigma^2 + \sigma^4 / 4) \omega_{t+1} + (\hat{\sigma}^2 + \hat{\sigma}^4 / 4) \hat{\omega}, \quad (5)$$

and

$$\mu_3\left(\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right) \middle| \omega_{t+1}\right) = -(3\sigma^4 / 2 + \sigma^6 / 8) \omega_{t+1} - (3\hat{\sigma}^4 / 2 + \hat{\sigma}^6 / 8) \hat{\omega}. \quad (6)$$

Equation (5) shows that the variance of the distribution of cross-sectional relative household consumption growth increases as household risk increases. Empirically, we find that the variance of cross-sectional relative household consumption growth is mildly countercyclical.

Equation (6) shows that the third central moment is always negative and becomes more so as household risk increases.⁷ Moreover, the model can generate any desired value of the third central moment, that is its maximum achievable absolute value is unbounded from above. Empirically, we find that the third central moment is mostly negative and mildly countercyclical. Guvenen, Ozkan, and Song (2014) find that the skewness of household income shocks is strongly countercyclical. This evidence allows us to associate an increase in household risk with recessions.⁸

For computational convenience, we define x_t in terms of the state variable ω_t as $x_t \equiv \left(e^{\gamma(\gamma-1)\sigma^2/2} - 1 \right) \omega_t$. Since x_t is proportional to ω_t , we sometimes refer to x_t as household risk, in place of ω_t . In our estimation, we limit the range of the RRA coefficient to $\gamma > 1$, which implies that $x_t > 0$. We assume that x_t follows an autoregressive gamma process, ARG (1), which is often referred to as a positive AR (1) process:

$$x_{t+1} = \nu\xi + \rho x_t + \varepsilon_{x,t+1}, \quad (7)$$

where $\nu > 0$, $\xi > 0$, and $\rho > 0$, $\varepsilon_{x,t+1}$ is a martingale difference sequence, $E[x_{t+1} | x_t] = \nu\xi + \rho x_t$, and $\text{var}(x_{t+1} | x_t) = \nu\xi^2 + 2\rho\xi x_t$.⁹ Thus, the autocorrelation of household risk is ρ . As we show below, the risk-free rate, price-dividend ratio, and expected market return are affine functions of household risk and therefore their autocorrelation is also ρ .

The heteroskedasticity of the innovation in household risk implies that the volatility of household risk is increasing in household risk, $var(x_{t+1} | x_t) = v\xi^2 + 2\rho\xi x_t$. This property drives key features of the economy. As we shall see shortly, the model implies that the variance of the risk-free rate, price-dividend ratio of the stock market, and expected market return are increasing in household risk and therefore are increasing in recessions.

In Appendix D, equation (D4), we calculate households' common SDF as

$$(SDF)_{t+1} = e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) - \gamma \Delta C_{t+1} + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} + \lambda x_{t+1}}, \quad (8)$$

where the parameters A_0 , A_1 , and λ are defined in Appendix D by equations (D2), (D3), and (D5).

The log risk-free rate is calculated in Appendix D, equation (D6), as

$$r_t = -\theta \log \delta - \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) - (\theta-1) (h_0 + h_1 A_0 - A_0) + \gamma \mu - \gamma^2 \sigma_a^2 / 2 - \lambda v \xi - \lambda^2 v \xi^2 / 2 - \{ \lambda \rho + \lambda^2 \rho \xi - (\theta-1) A_1 \} x_t. \quad (9)$$

Therefore, when household risk is high, the conditional variance of the risk-free rate is high. The model also implies that the risk-free rate is low when household risk is high since, in the estimated model, the coefficient on x_t in equation (9) is negative. Thus, the model implies that, in recessions, the risk-free rate is low and the variance of the risk-free rate is high. Both of these implications are consistent with observation.

Finally, the unconditional mean of the risk-free rate is

$$\begin{aligned} \bar{r} = & -\theta \log \delta - \hat{\omega} \left(e^{\gamma(\theta+1)\hat{\sigma}^2/2} - 1 \right) - (\theta-1)(h_0 + h_1 A_0 - A_0) + \gamma\mu - \gamma^2 \sigma_a^2 / 2 - \lambda v \xi - \lambda^2 v \xi^2 / 2 \\ & - \{ \lambda \rho + \lambda^2 \rho \xi - (\theta-1) A_1 \} \bar{x} \end{aligned} \quad (10)$$

and its unconditional variance is

$$\text{var}(r_t) = \{ \lambda \rho + \lambda^2 \rho \xi - (\theta-1) A_1 \}^2 \text{var}(x_t). \quad (11)$$

In Appendix D, we also show that the real yield curve is upward-sloping, downward-sloping, or humped depending on the state. Thus, the cross-sectional variation of the idiosyncratic income shocks gives rise to familiar shapes of the yield curve.

We assume that the log dividend growth of the stock market follows the process¹⁰

$$\Delta d_{t+1} = \alpha_d + \beta_d x_t + \sigma_d \varepsilon_{d,t+1}, \quad (12)$$

where $\varepsilon_{d,t+1} \sim \mathcal{N}(0,1)$ is i.i.d. and independent of all primitive random variables. The model yields a positive equity premium because the estimated value of the coefficient β_d is negative, capturing the scenario in which the expected dividend growth is decreasing in household risk (recessions).

In Appendix D, we calculate the price-dividend ratio in equation (D8) as

$$z_{m,t} = B_0 + B_1 x_t, \quad (13)$$

the expected stock market return in equation (D11) as

$$E[r_{m,t+1} | x_t] = k_0 + k_1 B_0 + k_1 B_1 v \xi - B_0 + \alpha_d + \beta_d x_t + \{k_1 B_1 \rho - B_1\} x_t, \quad (14)$$

and the conditional variance of the stock market return in equation (D12) as

$$\text{var}(r_{m,t+1} | x_t) = k_1^2 B_1^2 \text{var}(x_{t+1} | x_t) + \sigma_d^2, \quad (15)$$

where the parameters B_0 and B_1 are determined in Appendix D.

In the estimated model, the coefficient on x_t in equation (13) is negative, implying that the price-dividend ratio of the stock market is low in recessions. The coefficient on x_t in equation (14) is positive, implying that the expected return of the stock market is high in recessions. Finally, equation (15) implies that the variance of the stock market is high in recessions. All of these implications are consistent with the data.

II. Data Description

A. Prices and Dividends

We use monthly data on prices and dividends from January 1929 to December 2009. The proxy for the stock market is the Centre for Research in Security Prices (CRSP) value-weighted index of all stocks on the NYSE, AMEX, and NASDAQ. The monthly portfolio return is the sum of the portfolio price and dividend at the end of the month, divided by the portfolio price at the beginning of the month. The annual portfolio return is the sum of the portfolio price at the end of the year and uncompounded dividends over the year, divided by the portfolio price at the beginning of the year. The real annual portfolio return is the annual portfolio return deflated by the realized growth in the Consumer Price Index.

The proxy for the real annual risk-free rate is obtained as in Beeler and Campbell (2012). Specifically, the quarterly nominal yield on three-month Treasury bills is deflated using the realized growth in the Consumer Price Index to obtain the ex-post real three-month T-bill rate. The ex-ante quarterly risk-free rate is then obtained as the fitted value from the regression of the ex-post real three-month T-bill rate on the three-month nominal yield and the realized growth in the Consumer Price Index over the previous year. This procedure is equivalent to forecasting inflation and subtracting the inflation forecast from the three-month nominal yield. Finally, the ex-ante quarterly risk-free rate at the beginning of the year is annualized to obtain the ex-ante annual risk-free rate.

The above procedure in building an empirical counterpart of the real risk-free rate ensures that the real risk-free rate for a given period (say, a quarter) is known at the beginning of

the period, consistent with its definition in the context of our theoretical model. However, since the real risk-free rate is unobservable in the data over the duration of our sample period, we also consider two alternative approaches to constructing the real risk-free rate. In the first approach, we deflate the three-month nominal Treasury bill rate with realized inflation. The mean and volatility of the resulting time series are very similar to those obtained using our baseline procedure. In particular, for the entire 1947:Q1 to 2009:Q4 period over which aggregate quarterly consumption data are available, the mean of the risk-free rate is 0.003 and its volatility is 0.010. These values are very close (0.003 and 0.006, respectively) to those obtained using our baseline procedure. Similar conclusions obtain for the 1982:Q1 to 2009:Q4 subsample over which household consumption data are available: deflating the nominal three-month rate using realized inflation gives a mean risk-free rate of 0.006 with volatility 0.009 which are very close to the corresponding baseline values of 0.005 and 0.005, respectively.

In the second approach, we compound the monthly returns on one-month Treasury bills within a quarter to obtain the quarterly nominal risk-free rate and then convert it to real terms using realized inflation. This procedure also gives very similar values for the mean and volatility of the real risk-free rate: the mean and volatility are 0.002 and 0.009, respectively, over the 1947:Q1 to 2009:Q4 period and 0.005 and 0.009, respectively, over the 1982:Q1 to 2009:Q4 subperiod.

The above results suggest that our conclusions regarding the fit of the theoretical model for the risk-free rate is not sensitive to the precise procedure used to generate an empirical counterpart of the real risk-free rate.

The annual price-dividend ratio of the market is the market price at the end of the year, divided by the sum of dividends over the previous 12 months. The dividend growth rate is the

sum of dividends over the year divided by the sum of dividends over the previous year and is deflated using the realized growth in the Consumer Price Index.

*B. Household Consumption Data*¹¹

The household-level quarterly consumption data come from the Consumer Expenditure Survey (CEX) produced by the Bureau of Labor Statistics (BLS). This series of cross-sections covers the period since 1980:Q1. Each quarter, roughly 5,000 U.S. households are surveyed, chosen randomly according to stratification criteria determined by the U.S. Census. Each household participates in the survey for five consecutive quarters, one training quarter and four regular ones, during which time their recent consumption and other information are recorded. At the end of the fifth quarter, the subject household is replaced by another household, chosen randomly according to the stratification criteria determined by the U.S. Census. The cycle of subject households is staggered uniformly across quarters, such that new households replace approximately one-fifth of the participating households each quarter.^{12,13} If a household moves away from the sample address, it is dropped from the survey and the new household that moves into the address is included in the survey, provided it meets eligibility criteria.

The number of households in the database varies from quarter to quarter. The survey attempts to account for 95% of all quarterly household expenditures in each consumption category from a highly disaggregated list of consumption goods and services. At the end of the fourth regular quarter, data are also collected on the demographics and financial profiles of the participating households, including the value of asset holdings as of the month preceding the interview. We use consumption data only from the regular quarters, as we consider data from the

training quarter unreliable. In a significant number of years, the BLS failed to survey households not located near an urban area. We therefore limit attention to urban households.

The CEX survey reports are categorized into three tranches, namely, January, February, and March tranches. For a given year, first-quarter consumption for the January tranche corresponds to consumption over January through March, first-quarter consumption for the February tranche corresponds to consumption over February through April, first-quarter consumption for the March tranche corresponds to consumption over March through May and so on for second-, third-, and fourth-quarter consumption. While the CEX consumption data are presented at the monthly frequency for some consumption categories, the numbers reported as monthly are often simply the quarterly estimates divided by three.¹⁴ Thus, using monthly consumption is not an option.

Following Attanasio and Weber (1995), we discard from our sample consumption data for the years 1980 and 1981 because they are of questionable quality. Starting in interview period 1986:Q1, the BLS changed its household identification numbering system without providing the correspondence between the 1985:Q4 and 1986:Q1 identification numbers of households interviewed in both quarters. This change in the identification system makes it impossible to match households across the 1985:Q4 to 1986:Q1 gap and results in the loss of some observations. This problem recurs between 1996:Q1 and 1997:Q1 and also in 2005:Q1.

C. Definition of the Household Consumption Variables

For each tranche, we calculate each household's quarterly nondurables and services (NDS) consumption by aggregating the household's quarterly consumption across the

consumption categories that comprise NDS. We use consumption categories that adhere to the National Income and Product Accounts (NIPA) classification of NDS consumption. Since the quantity of interest here is relative household consumption growth, $\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right)$, it is unnecessary to either deflate or seasonally adjust consumption.

The per capita consumption, C_t , of a set of households is calculated as follows. First, the total consumption in a given quarter is obtained by summing the NDS consumption of all households in that quarter. Second, per capita consumption in a given quarter is obtained by dividing total consumption in that quarter by the sum of the number of family members across all households in that quarter. The per capita consumption growth between quarters $t - 1$ and t is defined as the ratio of the per capita consumption in quarters t and $t - 1$.

D. Household Selection Criteria

In each quarter, we delete from the sample those households that report zero total consumption, or zero consumption of NDS, or zero food consumption. We also delete household-quarters with missing information on the above items.

To mitigate observation error, we subject sample households to a consumption growth filter. The filter consists of the following selection criteria. First, we delete households with consumption growth reported in fewer than three consecutive quarters. Second, we delete the consumption growth rates $C_{i,t}/C_{i,t-1}$ and $C_{i,t+1}/C_{i,t}$ if $C_{i,t}/C_{i,t-1} < 1/2$ and $C_{i,t+1}/C_{i,t} > 2$, or if $C_{i,t}/C_{i,t-1} > 2$ and $C_{i,t+1}/C_{i,t} < 1/2$. Third, we delete consumption growth $C_{i,t}/C_{i,t-1}$ if it is greater than five.

III. Household Consumption Statistics

Using panel data on household consumption growth constructed from the CEX database, we compute summary statistics of the cross-sectional distribution of household consumption growth relative to per capita aggregate consumption growth, $\log\left(\frac{C_{i,t+1}/C_{t+1}}{C_{i,t}/C_t}\right)$. These statistics are reported in Table I for the period 1982:Q1 to 2009:Q4. The first row presents results for the January tranche. The second row presents results for households in the combined January, February, and March tranches, which provides a larger cross-section. The number of households that survive the filters described in Section II is substantially smaller than the original sample. For the January tranche, the maximum number of households in a quarter is 1,310 with a mean of 674.¹⁵ For the combined tranches, the maximum number of households in a quarter is 3,906 with a mean of 2,056.

-Table I here-

In the first row we find that, for the January tranche, the sample volatility, $\mu_2^{1/2}$, has a high mean value of 0.379 and is highly autocorrelated with a first-order autocorrelation coefficient of 0.772. The sample third central moment, μ_3 , is mostly negative with a mean value of -0.025 which is strongly statistically significant, and is positively autocorrelated with a first-order autocorrelation coefficient of 0.124. Our finding of negative skewness in the cross-sectional distribution of household consumption growth is consistent with that in Guvenen, Ozkan, and Song (2014), who report negative left skewness in the cross-sectional distribution of income growth.

While the third central moment is a commonly used measure of a distribution's skewness, we also present results for a quantile-based measure that is more robust to outliers. In particular, we report the time-series average of the fifth percentile of the relative household consumption growth distribution, along with its first-order autocorrelation coefficient. The average fifth percentile is equal to -0.595 and is strongly positively auto-correlated with a first-order autocorrelation coefficient of 0.534. The results are largely similar for the February and March tranches, which are omitted in the interest of brevity.

Stronger results obtain in the second row where we combine the households in the January, February, and March tranches to provide a larger cross-section. The third central moment is negative on average, strongly statistically significant, and positively autocorrelated. The fifth percentile of the cross-sectional distribution is negative and strongly positively autocorrelated with first-order autocorrelation coefficient of 0.792. In the lower panel of Table I, we present the statistics implied by the estimated model. We defer discussion of this panel until we present the empirical results.

In each quarter, the indicator variable I_{rec} takes the value of one if there is a NBER-designated recession in at least two of the three months in the quarter. In Table II, we present the correlations of the cross-sectional mean, volatility, third central moment, and fifth percentile with NBER-designated recessions. In recessions, the third central moment becomes more negative—the correlation between the recession dummy and the third central moment is -0.249 for the January tranche (row 1) and -0.217 for the combined January, February, and March tranches (row 2). The fifth percentile of the cross-sectional distribution also becomes more negative in recessions—the correlations between the recession dummy and the fifth percentile is -0.142 and -0.087, respectively, for the January tranche and the combined tranches. These

findings are also consistent with those reported in Guvenen, Ozkan, and Song (2014), who find evidence of countercyclical left skewness in the cross-sectional distribution of household income growth. In the last panel of Table II, we present correlations implied by the estimated model. We again defer discussion of this panel until we present the empirical results.

-Table II here-

Note that the CEX database may contain substantial measurement error because it is based on a survey. Measurement error may contaminate both the level of the variance and the third central moment of the cross-sectional distribution of household consumption growth, as well as their autocorrelations. To obtain more reliable estimates of households' consumption expenditures, Blundell, Pistaferri, and Preston (2008) impute a measure of NDS consumption for the families contained in the Panel Study of Income Dynamics (PSID) database by combining information from the PSID and CEX databases. They then use the imputed consumption to construct time series of the cross-sectional variance of household consumption growth for families in the PSID over the period 1980 to 1992. Furthermore, they specify a reduced-form model for the joint dynamics of household income and consumption that allows for measurement error in observed consumption. We apply their framework to compute the contribution of the different shocks, including measurement error, to the cross-sectional variance and third central moment of household consumption growth.

Blundell, Pistaferri, and Preston (2008) assume that the observed consumption of household i at time t may be written as the sum of the true unobserved consumption and measurement error that is independent over time:

$$c_{i,t}^* = c_{i,t} + u_{i,t}. \quad (16)$$

Therefore, taking first differences, we have

$$\Delta c_{i,t}^* = \Delta c_{i,t} + u_{i,t} - u_{i,t-1} = \underbrace{\phi \zeta_{i,t} + \psi \varepsilon_{i,t} + \xi_{i,t}}_{\Delta c_{i,t}} + u_{i,t} - u_{i,t-1}. \quad (17)$$

As in Blundell, Pistaferri, and Preston (2008), the true (unobserved) consumption growth, $\Delta c_{i,t}^*$, depends on permanent shocks to income (ζ), insurance with respect to the permanent shocks (ϕ), transitory shocks to income (ε), insurance with respect to the transitory shocks (ψ), and shocks to consumption that are uncorrelated with those to income (ξ), for example, preference shocks or shocks to the higher-order moments of the income process. All of the shocks are assumed to be mutually independent.

Given the above structure, we have

$$\text{var}(\Delta c_{i,t}^*) = \phi^2 \text{var}(\zeta_{i,t}) + \psi^2 \text{var}(\varepsilon_{i,t}) + \text{var}(\xi_{i,t}) + \text{var}(u_{i,t}) + \text{var}(u_{i,t-1}). \quad (18)$$

Note that using the CEX database, we estimate the time series of the cross-sectional volatility of relative household consumption growth (the square root of the left-hand side of (18)) over the 1982 to 2009 period. We find that, at the quarterly frequency, the cross-sectional volatility is 38% on average and varies from 28.4% to 80.5%.

Using the CEX and PSID databases to impute a measure of NDS consumption for the households in the PSID, Blundell, Pistaferri, and Preston (2008) show that, for the period 1980 to 1992, cross-sectional volatility is 39.6% on average and varies from 34.6% to 44.9%. These

estimates are similar to those obtained using the CEX data alone.

Based on the parameter estimates from Blundell, Pistaferri, and Preston (2008), the volatility of the cross-sectional household consumption growth accounted for by the permanent component (the square root of the first term on the right-hand side of (18)) is 8.7% on average, varying from 6.5% to 11.1%.¹⁶ The volatility of the cross-sectional household consumption growth accounted for by the permanent and transitory components (the square root of the sum of the first and second terms on the right-hand side of (18)) is 8.8% on average, varying from 6.6% to 11.2%. Thus, the transitory component contributes little to the overall volatility. Preference shocks and/or shocks to the higher-order moments of the income process (the third term on the right-hand side of (18)), together with the permanent and transitory shocks, generate an average cross-sectional volatility of 13.5% that varies from 12.2% to 15.2%. Therefore, measurement error (the fourth and fifth terms on the right-hand side of (18)) accounts for a large proportion of the overall cross-sectional volatility of household consumption growth.

Turning next to the third central moment, we have

$$E\left[\left(\Delta c_{i,t}^*\right)^3\right] = \phi^3 E\left(\zeta_{i,t}^3\right) + \psi^3 E\left(\varepsilon_{i,t}^3\right) + E\left(\zeta_{i,t}^3\right) + E\left(u_{i,t}^3\right) - E\left(u_{i,t-1}^3\right). \quad (19)$$

Assuming that $E\left(u_{i,t}^3\right) = E\left(u_{i,t-1}^3\right)$, equation (19) implies that measurement error does not affect the level of the third central moment.¹⁷

IV. Empirical Methodology and Results

A. Empirical Methodology

The model has 14 parameters: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the three parameters for household income shocks, σ , $\hat{\sigma}$, and $\hat{\omega}$; the three parameters for the dynamics of the state variable, ν , ξ , and ρ ; the three parameters governing the dynamics of aggregate dividend growth, α_b , β_d , and σ_d ; and the three preference parameters, namely the subjective discount factor, δ , the RRA coefficient, γ , and the EIS, ψ . We estimate the model parameters using GMM to simultaneously match the following moments of aggregate stock market data as well as household-level consumption data: the mean and variance of aggregate consumption growth rate; the mean, variance, and autocorrelation of the risk-free rate, market return, market-wide price-dividend ratio, and aggregate dividend growth rate; and the mean, variance, and third central moment of the cross-sectional distribution of household consumption growth relative to per capita aggregate consumption growth rate. We therefore have 17 moment restrictions in 14 parameters, providing overidentifying restrictions to test the model specification. We use a diagonal weighting matrix with a weight of one on all the moments except for the unconditional means of the market return and risk-free rate, which have weights of 100.¹⁸

B. Empirical Results

Since data on relative household consumption growth are available only at the quarterly frequency as of 1982:Q1, we estimate the model at the quarterly frequency over the period 1982:Q1 to 2009:Q4 to test the fit of the model-generated moments of aggregate stock market data and the cross-sectional distribution of relative quarterly household consumption growth to their sample counterparts. Table III presents model fit and parameter estimates for the January tranche.¹⁹ The table shows that the parsimonious model with just one state variable fits the sample moments of the risk-free rate, market return, price-dividend ratio, aggregate consumption growth rate, and dividend growth rate, as well as the first three central moments of the cross-sectional distribution of household consumption growth, very well. The *J*-stat is 4.38 and the model cannot be rejected at the 10% level of significance. The asymptotic 90% critical value is 6.52.

-Table III here-

The model generates a quarterly mean risk-free rate of 0.7% and stock market return of 2.0%, both very close to their sample counterparts of 0.5% and 1.9%, respectively. The model therefore provides an explanation for the equity premium and risk-free rate puzzles. The model generates volatility of 1.0% and a first-order autocorrelation of 0.876 for the risk-free rate, compared to sample counterparts of 0.5% and 0.863, respectively. The model also generates volatility of 22.8% and a first-order autocorrelation of -0.061 for the market return, compared to sample counterparts of 8.7% and 0.056, respectively. The model-implied mean of the market-wide price-dividend ratio is 3.787, which is very close to its sample counterpart of 3.745. More importantly, the model generates the high volatility of the price-dividend ratio observed in the

data (47.0% versus 41.1%), and thus explains the excess volatility puzzle. Note that most asset pricing models, including those with long-run risks and rare disasters, have difficulty matching the latter moment and therefore at explaining the high volatility of stock prices (see Constantinides and Ghosh (2011)). The model-implied first-order autocorrelation of the market-wide price-dividend ratio is 0.876, which is close to its sample counterpart of 0.972.

The model matches well the unconditional mean and volatility of the annual aggregate consumption growth rate. Note that models that rely on the incidence of shocks to aggregate, as opposed to household, consumption growth to address the equity premium and excess volatility puzzles require unrealistically high variance of aggregate consumption growth: the Barro (2006) rare disasters model implies aggregate consumption growth volatility of 4.6%. By contrast, the incidence of shocks to household consumption growth, as modeled in our paper, does not affect the volatility of aggregate consumption growth.

The model generates a mean of -0.2% and a volatility of 2.5% for the aggregate dividend growth rate, compared to sample counterparts of 0.5% and 2.6%, respectively. The model-implied autocorrelation of dividend growth rate is 0.308, close to its sample counterpart of 0.328.

The model matches well the unconditional third central moment of the cross-sectional distribution of relative household consumption growth. The model-implied third central moment of -0.027 is almost identical to its sample counterpart of -0.025. The model-implied unconditional volatility of the cross-sectional distribution of household consumption growth is 21.9%. Even though the volatility in the data is 37.9%, we argue in Section III that after adjusting for measurement error, the volatility is about 15%. Therefore, the model-implied unconditional volatility of the cross-sectional distribution of household consumption growth matches well its error-adjusted empirical counterpart.

The estimated preference parameters are reasonable: the risk aversion coefficient is 1.12 and the EIS is one. The EIS is higher than the inverse of the risk aversion coefficient, which highlights the importance of recursive preferences and points towards a preference for early resolution of uncertainty.

The parameters ν , ξ , and ρ govern household risk. The autocorrelation of household risk is $\rho=0.876$, which induces the autocorrelation of the risk-free rate and the price-dividend ratio to be 0.876 also, close to their sample values. The parameters ξ and ν govern the variance of household risk and induce the variance of the risk-free rate, market return, and price-dividend ratio to be close to their sample counterparts.

We use the point estimates of the model parameters in Table III to calculate the signs of the coefficients on household risk in the equations that determine the risk-free rate, price-dividend ratio, and conditional expected market return: $r_{f,t} = 0.009 - 28.38x_t$, $z_{m,t} = 3.88 - 1275.0x_t$, and $E[r_{m,t+1} | x_t] = 0.010 + 142.6x_t$. Consistent with empirical evidence, the model implies that the risk-free rate and price-dividend ratio are procyclical while the expected market return is countercyclical.

We extract the time series of the model-implied cross-sectional moments of household consumption growth from the observed time series of the risk-free rate and market-wide price-dividend ratio.²⁰ The middle panel of Table I displays the model-implied cross-sectional moments of household consumption growth. The volatility is of the same order of magnitude as its sample counterpart. The third central moment is almost identical to its sample counterpart. The first-order autocorrelation of the model-implied volatility is high and of the same order of magnitude as the autocorrelation in the data but the first-order autocorrelation of the third central moment is higher than the autocorrelation in the data, probably due to the small sample size and

the quality of the consumption data. The model-implied autocorrelation of the fifth percentile of the cross-sectional distribution is 84.6%, which is more in line with its sample counterpart of 53.4%. The bottom panel of Table I displays correlations between the historical and model-implied moments. The model-generated cross-sectional volatility has a correlation of 49.0% with its sample counterpart and the cross-sectional third central moment has a correlation of 37.7% with its sample counterpart.

The bottom panel of Table II displays the correlations between the model-implied household consumption growth moments and NBER-designated recessions. The correlation between the model-implied volatility of the cross-sectional distribution and recessions is 26.4%, the correlation between the third central moment and recessions is -26.8%, and the correlation between the fifth percentile and recessions is -23.4%. These values are close to the sample values of 13.9%, -24.9%, and -14.2%, respectively, for the January tranche and 11.8%, -21.7%, and -8.7%, respectively, for the combined tranches.

Note that we estimate the model without attempting to match the real term structure of interest rates. This is because data on long-term real risk-free bonds are only available for the post-2003 period in the U.S. The average annualized yield of 10-year real bonds, obtained from the Treasury website, is 0.019 over the period 2003 to 2009 (the end of our sample period). Based on the point estimates of the parameters, the model-implied annualized yield of 10-year real bonds is -0.024. While the term premium is slightly positive in the U.S. over the short available sample period, Evans (1998) documents that the premiums for inflation-indexed bonds in the U.K. are significantly negative (less than -2% at the one-year horizon). Given the short length of the period over which Treasury data are available, we do not attempt to determine

whether the real term structure is upward- or downward-sloping and we do not attempt to match such a pattern.

V. Further Interpretation of the Results

We consider several variants of the baseline model and obtain insights into the role each of its key ingredient plays. In Section V.A, we estimate the model over the entire quarterly sample period, 1947:Q1 to 2009:Q4, over which aggregate consumption and asset price data are available. In Section V.B, we examine the performance of the model at the annual frequency over the full sample period 1929 to 2009. In Section V.C, we investigate the role of recursive preferences by considering a variant of the baseline model in which households have constant RRA (CRRA) preferences. In Section V.D, we consider a version of the model in which the shocks to household consumption growth are conditionally lognormal, instead of a Poisson mixture of normal distributions as assumed in Section II. Finally, in Section V.E we highlight the significance of shocks unrelated to the business cycle by estimating a version of the model in which we suppress household income shocks that are unrelated to the business cycle.

A. Estimation with Quarterly Data, 1947:Q1 to 2009:Q4

We reestimate the model using quarterly data over the period 1947:Q1 to 2009:Q4, the period over which quarterly aggregate consumption data are available. Since household consumption data are not available over much of this period, the GMM system consists of only the 13 moment restrictions on aggregate asset prices and consumption and dividend growth rates.

The model fit and parameter estimates are presented in Table IV. The model matches well the moments of the risk-free rate, stock market return, and price-dividend ratio.

-Table IV here-

As with the quarterly analysis in Section IV, the model generates the empirically observed dynamics of the risk-free rate, price-dividend ratio, and stock market return. The model implies that the volatility of the risk-free rate, price-dividend ratio, and conditional expected market return are countercyclical, consistent with observation. We use the point estimates of the model parameters in Table IV to calculate the coefficients on household risk in the equations that determine the risk-free rate, price-dividend ratio, and conditional expected market return: $r_{f,t} = 0.010 - 0.188x_t$, $z_{m,t} = 3.73 - 31.07x_t$, and $E[r_{m,t+1} | x_t] = 0.010 + 0.768x_t$. Consistent with observation, the model implies that the risk-free rate and price-dividend ratio are procyclical while the expected market return is countercyclical. The estimated preference parameters are reasonable: the risk aversion coefficient is 13.3 and the EIS is 1.02. Based on the point estimates of the model parameters, the model-implied unconditional 10-year annualized real bond yield is 0.020, which is almost identical to its data counterpart of 0.019.

B. Estimation with Annual Data, 1929 to 2009

So far we have assumed that the decision frequency of households is quarterly. In this subsection, we present results at the annual frequency for the full sample period, 1929 to 2009. As in Section V.A, the GMM system consists of only the 13 moment restrictions on aggregate asset prices and aggregate consumption and dividend growth rates.

The results are reported in Table V. The model matches well the moments of the stock market return and price-dividend ratio but less so the moments of the risk-free rate. The model generates the empirically observed dynamics of the risk-free rate, price-dividend ratio, and stock market return. The model implies that the volatility of the risk-free rate, price-dividend ratio, and conditional expected market return are countercyclical, that the risk-free rate and price-dividend ratio are procyclical, and that the expected market return is countercyclical. The estimated preference parameters are reasonable: the RRA is 5.05 and the EIS is 1.10. These results suggest that the model is robust to the assumed decision frequency of households. Based on the point estimates of the model parameters, the model-implied unconditional 10-year real bond yield is 0.018, which is almost identical to its data counterpart of 0.019.

-Table V here-

C. Estimation with CRRA Preferences

To shed light on the role played by recursive preferences, we discuss the special case of our baseline model with CRRA preferences. In Table VI we present the counterpart of Tables III at the quarterly frequency but with CRRA preferences. Not surprisingly, we find that the equity premium and risk-free rate puzzles arise: the model-implied annual risk-free rate is 9.2% and the equity premium is 0.0%. In Table VII, we present the counterpart of Table V at the annual frequency but with CRRA preferences. Again we find that the equity premium and risk-free rate puzzles arise: the model-implied annual risk-free rate is 4.2% and the equity premium is -0.6%.

-Tables VI and VII here-

D. Estimation with Lognormal Shocks

In this section we present results for the same specification of preferences as in our baseline model but with idiosyncratic shocks to household consumption that are conditionally lognormal rather than negatively skewed:

$$\delta_{i,t} = \exp \left[\sum_{s=1}^t \left\{ (\sigma_s \eta_{i,s} - \sigma_s^2 / 2) + (\hat{\sigma} \hat{\eta}_{i,s} - \hat{\sigma}^2 / 2) \right\} \right], \quad \eta_{i,s}, \hat{\eta}_{i,s} \sim \text{i.i.d. } N(0,1). \quad (20)$$

The state variable σ_t^2 is assumed to follow an autoregressive gamma (positive AR (1)) process:

$$\sigma_{t+1}^2 = \nu_\sigma \xi_\sigma + \rho_\sigma \sigma_t^2 + \varepsilon_{\sigma,t+1}. \quad (21)$$

The results are reported in Table VIII for the quarterly subsample 1982:Q1 to 2009:Q4 for the January tranche. We find that the equity premium and risk-free rate puzzles arise: the model-implied annual risk-free rate is 8.0% and the equity premium is 0.4%, compared to sample counterparts of 2.0% and 6.4%, respectively.

-Table VIII here-

More importantly, the model implies that the equilibrium risk-free rate, price-dividend ratio, and conditional expected market return are affine functions of the state variable σ_t^2 . We

use the point estimates of the model parameters in Table VIII to calculate the coefficients on σ_t^2 in the equations that determine the risk-free rate, price-dividend ratio, and conditional expected market return: $r_{f,t} = 0.020 - 1.92\sigma_t^2$, $z_{m,t} = 3.82 - 141.1\sigma_t^2$, and $E[r_{m,t+1} | \sigma_t^2] = 0.020 + 6.70\sigma_t^2$. Note that the state variable σ_t^2 now also drives the variance of the cross-sectional distribution of household consumption growth. As a result, to generate the empirically observed dynamics of the risk-free rate, price-dividend ratio, and stock market return, the model requires that the variance of cross-sectional household consumption growth be countercyclical. This is improbable, given that Guvenen, Ozkan, and Song (2014) show that the variance of cross-sectional household income growth is not countercyclical.

E. Estimation without Household Income Shocks Unrelated to the Business Cycle

In this section we explore the significance of shocks to household income that are unrelated to the business cycle, by suppressing the term $\hat{j}_{i,s}^{1/2} \hat{\sigma} \hat{\eta}_{i,s} - \hat{j}_{i,s} \hat{\sigma}^2 / 2$ in equation (1) and writing shocks to household income as

$$\delta_{i,t} = \exp \left[\sum_{s=1}^t \left\{ \left(\hat{j}_{i,s}^{1/2} \hat{\sigma} \hat{\eta}_{i,s} - \hat{j}_{i,s} \hat{\sigma}^2 / 2 \right) \right\} \right]. \quad (22)$$

The results are reported in Table IX for the quarterly subsample 1982:Q1 to 2009:Q4 for the January tranche. The model-implied risk-free rate is negative and its volatility is too high. These

results are inferior to those in Table III and thus highlight the significance of shocks to household income that are unrelated to the business cycle.

-Table IX here-

VI. Household Consumption Risk and the Cross-Section of Excess Returns

Our empirical results above show that household consumption risk, measured by the third central moment of the cross-sectional distribution of household consumption growth, is an important risk factor that drives the time-series properties of aggregate quantities: the risk-free rate, market return, and market price-dividend ratio. We next show that household consumption risk also explains the cross-section of excess returns.

We follow the standard Fama-Macbeth (1973) methodology. In the first step, we run time-series regressions of quarterly excess returns of each asset on household consumption risk and obtain the factor loading for each asset. In the second step, for each quarter in the second half of the sample, we estimate a cross-sectional regression of excess asset returns on their estimated factor loadings from the first step and obtain a time series of cross-sectional intercepts and slope coefficients. We present the average of the cross-sectional intercepts, $\hat{\alpha}$, and slope coefficients, $\hat{\lambda}$. We calculate the standard errors of $\hat{\alpha}$ and $\hat{\lambda}$ from the time series of the cross-sectional intercepts and slope coefficients. Given the short length of the time series, we expect and find that the standard errors are large.

We present results for two variations of the first-stage time-series regressions. In the first variation (“rolling”), in each period t (starting with the midpoint of the sample) we use all of the returns up to period t to estimate the factor loadings as inputs to the cross-sectional regressions.

In the second variation (“fixed”), we use the first half of the sample to estimate the factor loadings as inputs to the cross-sectional regressions performed on the second half of the sample.

-Table X here-

The results for the rolling time-series regressions are reported in Table X. Panels A, B, and C present results when the set of test assets consists of the 25 size- and book-to-market-sorted equity portfolios of Fama and French (FF) (1993), the 30 industry-sorted portfolios, and the combined 25 FF and 30 industry-sorted portfolios, respectively. We include the industry portfolios as test assets in addition to the 25 FF portfolios because the size- and book-to-market-sorted equity portfolios have a strong factor structure making it easy for almost any proposed factor to produce a high cross-sectional adjusted R^2 (which we denote throughout by $\overline{R^2}$).²¹

In the first row of each panel, we present the results when the only factor is household consumption risk, the third central moment of the cross-sectional distribution of household consumption growth. In all three panels, the intercept is both statistically and economically insignificant, as expected. The slope coefficient is positive, as expected, but is not statistically significant given the small sample size. The cross-sectional $\overline{R^2}$ is stable, varying from 13.6% to 14.9%.

In the second row of each panel, we present the results when the only factor is the volatility of the cross-sectional distribution of household consumption growth. In all three panels, the intercept is both statistically and economically insignificant, as expected. In Panels A and C, the slope coefficient is negative, as expected, but small; in Panel B the slope coefficient is zero. The cross-sectional $\overline{R^2}$ varies from -6.9% to 40%, suggesting that the results are unstable and possibly spurious. Further evidence against volatility as a factor is provided in the third row of Panels A, B, and C, where we simultaneously consider household consumption risk and

volatility as factors. While in all three panels the slope coefficient on household consumption risk is positive as expected, in Panels B and C the slope on the volatility factor is positive, against expectation.

In the last row of each panel, we present the results for the three FF risk factors. In all three panels the estimated intercept is economically large; it is also statistically significant in Panel A. All slope coefficients are economically insignificant. The cross-sectional $\overline{R^2}$ varies from -22.8% to 59.5%, suggesting that the results are unstable.

The results for the fixed time-series regressions are reported in Table XI and reinforce the above results. When the only factor is household consumption risk, the intercept is both statistically and economically insignificant, as expected. The slope coefficient is positive, as expected, but is not statistically significant given the small sample size. The cross-sectional $\overline{R^2}$ is stable, varying from 7.5% to 21.5%.

-Table XI here-

When the only factor is the volatility of the cross-sectional distribution of household consumption growth, the intercept is both statistically and economically insignificant, as expected. The slope coefficient is positive in Panel B, against expectation and zero in Panel C, also against expectation. The cross-sectional $\overline{R^2}$ varies from -2.0% to 42.8%, suggesting that the results are unstable and possibly spurious.

With the three FF risk factors, the estimated intercept is economically large; it is also statistically significant in Panels B and C. All slope coefficients are economically insignificant. The cross-sectional $\overline{R^2}$ varies from 28.3% to 53.6%.

Overall, we conclude that household consumption risk does well in explaining the cross-section of excess returns: the intercept is economically and statistically insignificant, the slope

coefficient is consistently positive, as expected, and the cross-sectional adjusted $\overline{R^2}$ is consistently positive.

VII. Concluding Remarks

In this paper we explore cross-sectional variation in household income shocks as a factor that drives the time-series properties of the risk-free rate, market return, and market price-dividend ratio as well as the cross-section of excess returns. We focus on this channel by suppressing potential predictability of aggregate consumption growth and modeling it as an i.i.d. process. The model is parsimonious with only one state variable that is countercyclical and drives the moments of the cross-sectional distribution of household consumption growth. Despite this enforced parsimony, the model fits reasonably well both the unconditional and the conditional price moments, particularly the moments of the market price-dividend ratio, a target that has eluded a number of other models. More to the point, the model-generated moments of the cross-sectional distribution of household consumption match well their sample counterparts.

To check the robustness of model, we consider two variants of the baseline model. We first reestimate the model over the entire quarterly sample period, 1947:Q1 to 2009:Q4, over which aggregate consumption and asset price data are available and obtain similar results. We also reexamine the performance of the model at the annual frequency over the full sample period 1929 to 2009 and conclude that the results are insensitive to the assumed decision frequency of households.

We also explore the role played by three key ingredients of the model. First, we investigate the role of recursive preferences by considering a variant of the baseline model in

which households have CRRA preferences. We conclude that recursive preferences are pivotal in addressing the equity premium and risk-free rate puzzles. Second, we investigate the role of negative skewness in shocks to household consumption growth by considering a variant of the baseline model in which shocks to household consumption growth are conditionally lognormal. We conclude that negative skewness is pivotal in explaining the data. Third, we explore the significance of shocks to household income that are unrelated to the business cycle, in addition to shocks that are related to the business cycle. We find that shocks to household income that are unrelated to the business cycle also play a significant role.

Finally, we address the cross-sectional variation in the excess returns of size-, book-to-market-equity-, and industry-sorted equity portfolios. We show that the single factor—household consumption risk—explains the cross-section of excess returns at least as well as the combined three Fama-French factors do.

Appendix A: Proof that the Identity $I_t = C_t - D_t$ is Respected

Since the households are symmetric and their number is normalized to equal one, we apply the law of large numbers as in Green (1989) and claim that $I_t = E[I_{i,t} | C_t, D_t]$. Furthermore, since household shocks are assumed to be conditionally normally distributed and independent of anything else in the economy, we obtain

$$\begin{aligned}
 I_t &= E[I_{i,t} | C_t, D_t] \\
 &= E \left[\exp \left(\sum_{s=1}^t \left\{ j_{i,s}^{1/2} \sigma \eta_{i,s} - j_{i,s} \sigma^2 / 2 + \hat{j}_{i,s}^{1/2} \hat{\sigma} \hat{\eta}_{i,s} - \hat{j}_{i,s} \hat{\sigma}^2 / 2 \right\} \right) \right] C_t - D_t \\
 &= E \left[E \left[\exp \left(\left\{ \sum_{s=1}^t j_{i,s}^{1/2} \sigma \eta_{i,s} - j_{i,s} \sigma^2 / 2 + \hat{j}_{i,s}^{1/2} \hat{\sigma} \hat{\eta}_{i,s} - \hat{j}_{i,s} \hat{\sigma}^2 / 2 \right\} \right) \middle| \left\{ J_{i,\tau}, \hat{J}_{i,\tau} \right\}_{\tau=1,\dots,t} \right] \right] C_t - D_t \\
 &= C_t - D_t, \tag{A1}
 \end{aligned}$$

proving the claim.

Appendix B: Proof that Autarchy is an Equilibrium

We conjecture and verify that autarchy is an equilibrium. The proof follows several steps. First, we calculate household i 's private valuation of its wealth portfolio. Next we calculate the log return, $r_{i,c,t+1}$, on household i 's wealth portfolio and substitute this return in the household's SDF, as stated in equation (3). We integrate out of this SDF the household's idiosyncratic income shocks and show that households have a common SDF. This implies that the private valuation of any security with payoffs independent of the idiosyncratic income shocks is the same across households, verifying the conjecture that autarchy is an equilibrium.

Let $P_{i,c,t}$ be the price of household i 's private valuation of its wealth portfolio, $Z_{i,c,t} \equiv P_{i,c,t} / C_{i,t}$, and $z_{i,c,t} \equiv \log(Z_{i,c,t})$. We prove by induction that the price-to-consumption ratio is a function of only state variable ω_t . We conjecture that $z_{i,c,t+1} = z_{c,t+1}(\omega_{t+1})$. The Euler equation for $r_{i,c,t+1}$ is

$$E \left[e^{\frac{\theta \log \delta - \frac{\theta}{\psi} \Delta c_{i,t+1} + (\theta-1)r_{i,c,t+1} + r_{i,c,t+1}}{\psi}} \mid \Delta c_t, \omega_t, j_{i,t}, \eta_{i,t}, \hat{j}_{i,t}, \hat{\eta}_{i,t} \right] = 1. \quad (\text{B1})$$

We write

$$\begin{aligned} r_{i,c,t+1} &= \log(P_{i,c,t+1} + C_{i,t+1}) - \log P_{i,c,t} \\ &= \log(Z_{i,c,t+1} + 1) - \log(Z_{i,c,t}) + \log C_{i,t+1} - \log C_{i,t} \\ &= \log(e^{z_{c,t+1}} + 1) - z_{i,c,t} + \Delta c_{i,t+1} \end{aligned} \quad (\text{B2})$$

and substitute (B2) into Euler equation (B1) to obtain

$$E \left[e^{\theta \log \delta - \frac{\theta}{\psi} \Delta c_{i,t+1} + \theta \left(\log \left(e^{z_{i,c,t+1}} + 1 \right) - z_{i,c,t+1} \right)} \mid \Delta c_t, \omega_t, j_{i,t}, \eta_{i,t}, \hat{j}_{i,t}, \hat{\eta}_{i,t} \right] = 1$$

or

$$e^{\theta z_{i,c,t}} = E \left[e^{\theta \log \delta + (1-\gamma) \left(\mu + \sigma_a \varepsilon_{t+1} + j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) + \theta \log \left(e^{z_{c,t+1}} + 1 \right)} \mid \Delta c_t, \omega_t, j_{i,t}, \eta_{i,t}, \hat{j}_{i,t}, \hat{\eta}_{i,t} \right]. \quad (\text{B3})$$

We integrate out of equation (B3) the random variables $\varepsilon_{t+1}, \eta_{i,t+1}, j_{i,t+1}, \hat{\eta}_{i,t+1}$, and $\hat{j}_{i,t+1}$, leaving $z_{i,c,t}$ as a function of only ω_t , which proves the claim that $z_{i,c,t} = z_{c,t}(\omega_t)$.

Now, the $(SDF)_{i,t+1}$ of household i is

$$\begin{aligned} (SDF)_{i,t+1} &= \exp \left(\theta \log \delta - \frac{\theta}{\psi} \Delta c_{i,t+1} + (\theta - 1) r_{i,c,t+1} \right) \\ &= \exp \left(\theta \log \delta - \gamma \left(\Delta c_{t+1} + j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) + (\theta - 1) \left(\log \left(e^{z_{c,t+1}} + 1 \right) - z_{c,t} \right) \right). \end{aligned} \quad (\text{B4})$$

In pricing any security other than households' wealth portfolios, we integrate out of $(SDF)_{i,t+1}$ the household-specific random variables $\eta_{i,t+1}, j_{i,t+1}, \hat{\eta}_{i,t+1}$, and $\hat{j}_{i,t+1}$ and obtain a SDF common across households. Therefore, each household's private valuation of any security, other than households' wealth portfolios, is common. This completes the proof that no-trade is an equilibrium.

Appendix C: Derivation of the Cross-Sectional Moments of Consumption Growth

We use the following result:

$$e^{-\omega} \sum_{n=0}^{\infty} e^{kn} \omega^n / n! = e^{-\omega} \sum_{n=0}^{\infty} (e^k \omega)^n / n! = e^{-\omega} e^{e^k \omega}. \quad (\text{C1})$$

Differentiating one, two, and three times with respect to k and setting $k = 0$, we obtain

$$\begin{aligned} e^{-\omega} \sum_{n=0}^{\infty} n \omega^n / n! &= \omega \\ e^{-\omega} \sum_{n=0}^{\infty} n^2 \omega^n / n! &= \omega^2 + \omega \\ e^{-\omega} \sum_{n=0}^{\infty} n^3 \omega^n / n! &= \omega^3 + 3\omega^2 + \omega. \end{aligned} \quad (\text{C2})$$

We calculate the mean of the cross-sectional distribution of relative household consumption growth as

$$\begin{aligned}
\mu_1 &= E \left[\log \left(\frac{C_{i,t+1} / C_{t+1}}{C_{i,t} / C_t} \right) \middle| \omega_{t+1} \right] \\
&= E \left[E \left[\log \left(\frac{C_{i,t+1} / C_{t+1}}{C_{i,t} / C_t} \right) \middle| j_{i,t+1}, \hat{j}_{i,t+1} \right] \middle| \omega_{t+1} \right] \\
&= E \left[E \left[j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \middle| j_{i,t+1}, \hat{j}_{i,t+1} \right] \middle| \omega_{t+1} \right] \quad (C3) \\
&= E \left[-j_{i,t+1} \sigma^2 / 2 - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \middle| \omega_{t+1} \right] \\
&= -(\sigma^2 / 2) \omega_{t+1} - (\hat{\sigma}^2 / 2) \hat{\omega}.
\end{aligned}$$

We calculate the variance as

$$\begin{aligned}
\mu_2 &= \text{var} \left(\log \left(\frac{C_{i,t+1} / C_{t+1}}{C_{i,t} / C_t} \right) \middle| \omega_{t+1} \right) \\
&= \text{var} \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) \middle| \omega_{t+1} \right] \\
&= \text{var} \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 \right) \middle| \omega_{t+1} \right] + \text{var} \left[\left(\hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) \middle| \omega_{t+1} \right] \\
&= E \left[E \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 \right)^2 \middle| j_{i,t+1} \right] \middle| \omega_{t+1} \right] - (\sigma^2 \omega_{t+1} / 2)^2 \\
&\quad + E \left[E \left[\left(\hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right)^2 \middle| \hat{j}_{i,t+1} \right] \right] - (\hat{\sigma}^2 \hat{\omega} / 2)^2 \\
&= E \left[\left(j_{i,t+1} \sigma^2 + j_{i,t+1}^2 \sigma^4 / 4 \right) \middle| \omega_{t+1} \right] - (\sigma^2 \omega_{t+1} / 2)^2 + E \left[\hat{j}_{i,t+1} \hat{\sigma}^2 + \hat{j}_{i,t+1}^2 \hat{\sigma}^4 / 4 \right] - (\hat{\sigma}^2 \hat{\omega} / 2)^2 \\
&= \sigma^2 \omega_{t+1} + (\sigma^4 / 4) \omega_{t+1} (1 + \omega_{t+1}) - (\sigma^2 \omega_{t+1} / 2)^2 + \hat{\sigma}^2 \hat{\omega} + (\hat{\sigma}^4 / 4) \hat{\omega} (1 + \hat{\omega}) - (\hat{\sigma}^2 \hat{\omega} / 2)^2 \\
&= (\sigma^2 + \sigma^4 / 4) \omega_{t+1} + (\hat{\sigma}^2 + \hat{\sigma}^4 / 4) \hat{\omega}.
\end{aligned} \tag{C4}$$

We calculate the third central moment as

$$\begin{aligned}
& \mu_3 \left(\log \left(\frac{C_{i,t+1} / C_{t+1}}{C_{i,t} / C_t} \right) \middle| \omega_{t+1} \right) \\
&= \mu_3 \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) \middle| \omega_{t+1} \right] \\
&= \mu_3 \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 \right) \middle| \omega_{t+1} \right] + \mu_3 \left[\left(\hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) \middle| \omega_{t+1} \right].
\end{aligned}$$

We have

$$\begin{aligned}
& \mu_3 \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 \right) \middle| \omega_{t+1} \right] \\
&= E \left[E \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} + (\omega_{t+1} - j_{i,t+1}) \sigma^2 / 2 \right)^3 \middle| j_{i,t+1} \right] \middle| \omega_{t+1} \right] \\
&= \sigma^3 E \left[E \left[\left(j_{i,t+1}^{1/2} \eta_{i,t+1} + (\omega_{t+1} - j_{i,t+1}) \sigma / 2 \right)^3 \middle| j_{i,t+1} \right] \middle| \omega_{t+1} \right] \\
&= \sigma^3 E \left[E \left[3 j_{i,t+1} (\omega_{t+1} - j_{i,t+1}) \sigma / 2 + (\omega_{t+1} - j_{i,t+1})^3 \sigma^3 / 8 \middle| j_{i,t+1} \right] \middle| \omega_{t+1} \right] \\
&= (\sigma^4 / 2) E \left[E \left[3 j_{i,t+1} \omega_{t+1} - 3 j_{i,t+1}^2 + (\omega_{t+1}^3 - 3 \omega_{t+1}^2 j_{i,t+1} + 3 \omega_{t+1} j_{i,t+1}^2 - j_{i,t+1}^3) \sigma^2 / 4 \middle| j_{i,t+1} \right] \middle| \omega_{t+1} \right] \\
&= (\sigma^4 / 2) \left\{ 3 \omega_{t+1}^2 - 3 (\omega_{t+1}^2 + \omega_{t+1}) + (\omega_{t+1}^3 - 3 \omega_{t+1}^3 + 3 \omega_{t+1} (\omega_{t+1}^2 + \omega_{t+1})) - (\omega_{t+1}^3 + 3 \omega_{t+1}^2 + \omega_{t+1}) \right\} \sigma^2 / 4 \\
&= -(3 \sigma^4 / 2 + \sigma^6 / 8) \omega_{t+1}.
\end{aligned}$$

Likewise, we can show that $\mu_3 \left(\hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) = - \left(3 \hat{\sigma}^4 / 2 + \hat{\sigma}^6 / 8 \right) \hat{\omega}$. Therefore,

$$\mu_3 \left(\log \left(\frac{C_{i,t+1} / C_{t+1}}{C_{i,t} / C_t} \right) \middle| \omega_{t+1} \right) = -(3 \sigma^4 / 2 + \sigma^6 / 8) \omega_{t+1} - \left(3 \hat{\sigma}^4 / 2 + \hat{\sigma}^6 / 8 \right) \hat{\omega}. \quad (C5)$$

**Appendix D: Derivation of the Common SDF, Risk-Free Rate,
Market Price-Dividend Ratio, and Expected Market Return**

Solution for a Household's Consumption-Wealth Ratio. In Appendix B we prove that any household's consumption-wealth ratio is a function of only the state variable, that is, $z_{i,c,t} = z_{c,t}(\omega_t)$. Here we conjecture and verify that $z_{c,t} = A_0 + A_1 x_t$. To do so, we plug $z_{c,t} = A_0 + A_1 x_t$ into the Euler equation (B3). We also log-linearize the term $\log(e^{z_{c,t+1}} + 1)$ as in Campbell and Shiller (1988) and obtain $\log(e^{z_{c,t+1}} + 1) \approx h_0 + h_1 z_{c,t+1}$, where

$$h_0 \equiv \log\left(e^{\bar{z}_c} + 1\right) - \frac{\bar{z}_c e^{\bar{z}_c}}{e^{\bar{z}_c} + 1} \text{ and } h_1 \equiv \frac{e^{\bar{z}_c}}{e^{\bar{z}_c} + 1}. \text{ We have specifically,}$$

$$e^{\theta z_{i,c,t}} = E \left[e^{\theta \log \delta (1-\gamma) \left(\mu + \sigma_a \varepsilon_{i,t} + J_{i,t}^{1/2} \sigma_{\eta_{i,t}} \eta_{i,t} - J_{i,t} \sigma^2 / 2 + \hat{J}_{i,t}^{1/2} \hat{\sigma} \hat{\eta}_{i,t} - \hat{J}_{i,t} \hat{\sigma}^2 / 2 \right) + \theta \log \left(e^{z_{c,t+1}} + 1 \right)} \mid \Delta c_t, \omega_t, J_{i,t}, \eta_{i,t}, \hat{J}_{i,t}, \hat{\eta}_{i,t} \right]$$

or

$$e^{\theta(A_0 + A_1 x_t)} = E \left[e^{\theta \log \delta (1-\gamma) \left(\mu + \sigma_a \varepsilon_{i,t} + J_{i,t}^{1/2} \sigma_{\eta_{i,t}} \eta_{i,t} - J_{i,t} \sigma^2 / 2 + \hat{J}_{i,t}^{1/2} \hat{\sigma} \hat{\eta}_{i,t} - \hat{J}_{i,t} \hat{\sigma}^2 / 2 \right) + \theta \left(h_0 + h_1 (A_0 + A_1 x_{t+1}) \right)} \mid \Delta c_t, \omega_t, J_{i,t}, \eta_{i,t}, \hat{J}_{i,t}, \hat{\eta}_{i,t} \right]$$

or

$$E \left[e^{\theta \log \delta (1-\gamma) \mu (1-\gamma) \sigma_a^2 / 2 + \gamma (1-\gamma) \left(J_{i,t} + \hat{J}_{i,t} \right) + \theta \left(h_0 + h_1 (A_0 + A_1 x_{t+1}) - A_0 - A_1 x_t \right)} \mid \omega_t, J_{i,t}, \hat{J}_{i,t} \right] = 1$$

or

$$E \left[e^{\theta \log \delta + (1-\gamma)\mu + (1-\gamma)^2 \sigma_a^2 / 2 + x_{t+1} + \left(e^{\gamma(\gamma-1)\hat{\sigma}^2/2} - 1 \right) \hat{\omega} + \theta \{ h_0 + h_1 (A_0 + A_1 x_{t+1}) - A_0 - A_1 x_t \}} \mid \omega_t \right] = 1,$$

since $e^{-\omega} \sum_{n=0}^{\infty} e^{kn} \omega^n / n! = e^{-\omega} \sum_{n=0}^{\infty} (e^k \omega)^n / n! = e^{-\omega} e^{e^k \omega}$ and $\left(e^{(\gamma-1)\gamma\sigma^2/2} - 1 \right) \omega_t \equiv x_t$. Therefore,

$$e^{\theta \log \delta + (1-\gamma)\mu + (1-\gamma)^2 \sigma_a^2 / 2 + \left(e^{\gamma(\gamma-1)\hat{\sigma}^2/2} - 1 \right) \hat{\omega} + \theta (h_0 + h_1 A_0 - A_0 - A_1 x_t)} E \left[e^{(1+\theta h_1 A_1) x_{t+1}} \mid \omega_t \right] = 1$$

or, up to a second-order approximation,

$$e^{\theta \log \delta + (1-\gamma)\mu + (1-\gamma)^2 \sigma_a^2 / 2 + \left(e^{\gamma(\gamma-1)\hat{\sigma}^2/2} - 1 \right) \hat{\omega} + \theta (h_0 + h_1 A_0 - A_0 - A_1 x_t) + (1+\theta h_1 A_1)(v\xi + \rho x_t) + (1+\theta h_1 A_1)^2 (v\xi^2 + 2\rho\xi x_t) / 2} \approx 1. \quad (\text{D1})$$

Matching the constant, we obtain

$$\theta \log \delta + (1-\gamma)\mu + (1-\gamma)^2 \sigma_a^2 / 2 + \left(e^{\gamma(\gamma-1)\hat{\sigma}^2/2} - 1 \right) \hat{\omega} + \theta (h_0 + h_1 A_0 - A_0) + (1+\theta h_1 A_1) v\xi + \frac{1}{2} (1+\theta h_1 A_1)^2 v\xi^2 = 0, \quad (\text{D2})$$

and matching the coefficient of x_t , we obtain

$$-A_1\theta + (1 + \theta h_1 A_1)\rho + (1 + \theta h_1 A_1)^2 \rho \xi = 0. \quad (\text{D3})$$

The solution of equations (D2) and (D3) produces values for the parameters A_0 and A_1 that verify

the conjecture that $z_{c,t} = A_0 + A_1 x_t$.²² Since $\bar{z}_c = A_0 + A_1 \bar{x}$, $h_0 \equiv \log\left(e^{\bar{z}_c} + 1\right) - \frac{\bar{z}_c e^{\bar{z}_c}}{e^{\bar{z}_c} + 1}$, and

$h_1 \equiv \frac{e^{\bar{z}_c}}{e^{\bar{z}_c} + 1}$, the parameters h_0 and h_1 are determined in terms of the parameters A_0 , A_1 , and \bar{x} .

Common SDF across Households. In pricing any security, other than households' wealth portfolios, we integrate out of the SDF in equation (B4) the household-specific random variables

$\theta_{i,t+1}$, $j_{i,t+1}$, $\hat{\theta}_{i,t+1}$, and $\hat{j}_{i,t+1}$ and obtain a SDF that is common across households:

$$\begin{aligned} (SDF)_{t+1} &= E \left[e^{\theta \log \delta - \gamma \left(\Delta c_{t+1} + j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} - j_{i,t+1} \sigma^2 / 2 + \hat{j}_{i,t+1}^{1/2} \hat{\sigma} \hat{\eta}_{i,t+1} - \hat{j}_{i,t+1} \hat{\sigma}^2 / 2 \right) + (\theta - 1) (h_0 + h_1 z_{c,t+1} - z_{c,t})} \mid c_t, c_{t+1}, \omega_t, \omega_{t+1} \right] \\ &= e^{\theta \log \delta - \gamma \Delta c_{t+1} + \omega_{t+1} \left(e^{\gamma(\gamma+1)\sigma^2/2} - 1 \right) + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) + (\theta - 1) (h_0 + h_1 z_{c,t+1} - z_{c,t})} \\ &= e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) - \gamma \Delta c_{t+1} + (\theta - 1) (h_0 + h_1 A_0 - (A_0 + A_1 x_t)) + \lambda x_{t+1}}, \end{aligned} \quad (\text{D4})$$

where

$$\lambda \equiv \frac{e^{\gamma(\gamma+1)\sigma^2/2} - 1}{e^{\gamma(\gamma-1)\sigma^2/2} - 1} + (\theta - 1) h_1 A_1. \quad (\text{D5})$$

Solution for the Risk-Free Rate. The Euler equation for the log risk-free rate is

$$E \left[e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) - \gamma \Delta c_{t+1} + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} + \lambda x_{t+1} + r_t} \mid \omega_t \right] = 1,$$

or, up to a second-order approximation,

$$e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} - \gamma \mu + \gamma^2 \sigma_a^2 / 2 + \lambda (v \xi + \rho x_t) + \lambda^2 (v \xi^2 + 2 \rho \xi x_t) / 2 + r_t} \approx 1$$

or

$$r_t = -\theta \log \delta - \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) - (\theta-1) (h_0 + h_1 A_0 - A_0) + \gamma \mu - \gamma^2 \sigma_a^2 / 2 - \lambda v \xi - \lambda^2 v \xi^2 / 2 - \{ \lambda \rho + \lambda^2 \rho \xi - (\theta-1) A_1 \} x_t. \quad (D6)$$

The implications of the model regarding the term structure of interest rates are the same as those for a discretized version of the Cox, Ingersoll, and Ross (1985) model. Recall that x_t follows a heteroskedastic AR (1) process with conditional variance $v \xi^2 + 2 \rho \xi x_t$. We prove that, under the risk-neutral probability measure Q , x_t follows a heteroskedastic AR (1) process, where the mean of x_{t+1} , conditional on x_t , is shifted by approximately $\lambda [v \xi^2 + 2 \rho \xi x_t]$ and the variance is affine in x_t . To see this, note that $e^{r_t} (SDF)_{t+1}$ is the discrete-time Radon-Nikodym derivative. Under the risk-neutral probability measure Q , the mean of x_{t+1} , conditional on x_t , is

$E^Q[x_{t+1} | x_t] = E[x_{t+1} e^{r_t} (SDF)_{t+1} | x_t] \approx E[x_{t+1} | x_t] + \lambda [v\xi^2 + 2\rho\xi x_t]$ and its variance is affine in x_t . Since the risk-free rate is affine in household risk x_t , the risk-free rate also follows a heteroskedastic AR (1) process with variance of the innovation affine in the risk-free rate under the risk-neutral probability measure. Thus, the model is isomorphic to a discretized version of the Cox, Ingersoll, and Ross (1985) model and implies that the yield curve is upward-sloping, downward-sloping, or humped depending on the state, where the state may be represented by the risk-free rate.

The Price-Dividend Ratio of the Stock Market. We denote the log stock market return as $r_{m,t}$ and the stock market price-dividend ratio as $z_{m,t}$. As in Campbell-Shiller (1988), we write

$$r_{m,t+1} = k_0 + k_1 z_{m,t+1} - z_{m,t} + \Delta d_{t+1}, \quad (\text{D7})$$

where $k_0 \equiv \log(e^{\bar{z}_m} + 1) - \frac{\bar{z}_m e^{\bar{z}_m}}{e^{\bar{z}_m} + 1}$ and $k_1 = \frac{e^{\bar{z}_m}}{e^{\bar{z}_m} + 1}$. We conjecture and verify that the price-

dividend ratio of the stock market is

$$z_{m,t} = B_0 + B_1 x_t \quad (\text{D8})$$

and write

$$r_{m,t+1} = k_0 + k_1 (B_0 + B_1 x_{t+1}) - (B_0 + B_1 x_t) + \alpha_d + \beta_d x_t + \sigma_d \varepsilon_{d,t+1}.$$

The Euler equation is

$$E \left[e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\sigma_d^2/2} - 1 \right) + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} - \gamma \Delta c_{t+1} + \lambda x_{t+1} + r_{m,t+1}} \mid \omega_t \right] = 1$$

or

$$E \left[e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\sigma_d^2/2} - 1 \right) + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} - \gamma \Delta c_{t+1} + \lambda x_{t+1} + k_0 + k_1 (B_0 + B_1 x_{t+1}) - (B_0 + B_1 x_t) + \alpha_d + \beta_d x_t + \sigma_d \varepsilon_{d,t+1}} \mid \omega_t \right] = 1$$

or

$$e^{\theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\sigma_d^2/2} - 1 \right) + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} - \gamma \mu + \gamma^2 \sigma_d^2 / 2 + k_0 + k_1 B_0 - (B_0 + B_1 x_t) + \alpha_d + \beta_d x_t + \sigma_d^2 / 2} E \left[e^{(\lambda + k_1 B_1) x_{t+1}} \mid \omega_t \right] = 1,$$

or, up to a second-order approximation,

$$\begin{aligned} & \theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\sigma_d^2/2} - 1 \right) + (\theta-1) \{ h_0 + h_1 A_0 - (A_0 + A_1 x_t) \} - \gamma \mu + \gamma^2 \sigma_d^2 / 2 + k_0 + k_1 B_0 \\ & - (B_0 + B_1 x_t) + \alpha_d + \beta_d x_t + \sigma_d^2 / 2 + (\lambda + k_1 B_1) \{ \nu \xi + \rho x_t \} + (\lambda + k_1 B_1)^2 \{ \nu \xi^2 + 2 \rho \xi x_t \} / 2 \\ & \approx 0. \end{aligned}$$

We set the constant and coefficient on x_t equal to zero and obtain two equations that determine

the parameters B_0 and B_1 :

$$\begin{aligned}
& \theta \log \delta + \hat{\omega} \left(e^{\gamma(\gamma+1)\hat{\sigma}^2/2} - 1 \right) + (\theta - 1)(h_0 + h_1 A_0 - A_0) - \gamma \mu \\
& + \gamma^2 \sigma_a^2 / 2 + k_0 + k_1 B_0 - B_0 + \alpha_d + \sigma_d^2 / 2 + (\lambda + k_1 B_1) \nu \xi + \frac{1}{2} (\lambda + \kappa_1 B_1)^2 \nu \xi^2 \\
& = 0
\end{aligned} \tag{D9}$$

and

$$-(\theta - 1) A_1 - B_1 + \beta_d + (\lambda + k_1 B_1) \rho + (\lambda + k_1 B_1)^2 \rho \xi = 0. \tag{D10}$$

Note that the parameters k_0 and k_1 are determined in terms of the parameters B_0 , B_1 , and \bar{x} .

Therefore, the expected stock market return is

$$\begin{aligned}
E[r_{m,t+1} | \omega_t] &= k_0 + k_1 B_0 + k_1 B_1 \{ \nu \xi + \rho x_t \} - (B_0 + B_1 x_t) + \alpha_d + \beta_d x_t \\
&= k_0 + k_1 B_0 + k_1 B_1 \nu \xi - B_0 + \alpha_d + \beta_d x_t + \{ k_1 B_1 \rho - B_1 \} x_t.
\end{aligned} \tag{D11}$$

REFERENCES

- Ahn, Seung C., and Christopher S. Gadarowski, 1999, Small sample properties of the model specification test based on the Hansen-Jagannathan distance, working paper, Arizona State University.
- Attanasio, Orazio P., and Steven J. Davis, 1996, Relative wage movements and the distribution of consumption, *Journal of Political Economy* 104, 1227-1262.
- Attanasio, Orazio P., and Guliermo Weber, 1995, Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey, *Journal of Political Economy* 103, 1121-1157.
- Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters implied by equity index options, *Journal of Finance* 66, 1967-2009.
- Barro, Robert J., 2006, Rare disasters and asset markets in the 20th century, *Quarterly Journal of Economics* 121, 823-866.
- Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* 39, 255-335.
- Beeler, Jason, and John Y. Campbell, 2012, The long-run risks model and aggregate asset prices: An empirical assessment, *Critical Finance Review* 1, 141-182.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston, 2008, Consumption inequality and partial insurance, *American Economic Review* 98, 1887-1921.
- Brav, Alon, George M. Constantinides, and Christopher C. Geczy, 2002, Asset pricing with

heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793-824.

Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195-228.

Cochrane, John, 1991, A simple test of consumption insurance, *Journal of Political Economy* 99, 957-976.

Cogley, Thomas, 2002, Idiosyncratic risk and the equity premium: Evidence from the Consumer Expenditure Survey, *Journal of Monetary Economics* 49, 309-334.

Constantinides, George M., 1982, Intertemporal asset pricing with heterogeneous consumers and without demand aggregation, *Journal of Business* 55, 253-267.

Constantinides, George M., 2008, Comment on Barro and Ursua, *Brookings Papers on Economic Activity* 39, 341-350.

Constantinides, George M., and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219-240.

Constantinides, George M., and Anisha Ghosh, 2011, Asset pricing tests with long-run risks in consumption growth, *Review of Asset Pricing Studies* 1, 96-136.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385-407.

Dew-Becker, Ian, and Stefano Giglio, 2016, Asset pricing in the frequency domain: theory and empirics, *Review of Financial Studies* 29, 2029-2068.

Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki, 2014, How much would you pay to resolve long-run risk? *American Economic Review* 104, 2680-2697.

Epstein, Larry G., and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Evans, Martin D. D., 1998, Real rates, expected inflation and inflation risk premia, *Journal of Finance* 53, 187-218.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns of stocks and bonds, *Journal of Financial Economics* 33, 3-56.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: empirical tests, *Journal of Political Economy* 81, 607-636.

Ferson, Wayne, and Stephen R. Foerster, 1994, Finite sample properties of the generalized methods of moments tests of conditional asset pricing models, *Journal of Financial Economics* 36, 29-56.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645-700.

Ghosh, Anisha, Christian Julliard, and Alex P. Taylor, 2014, What is the consumption-CAPM missing? An information-theoretic framework for the analysis of asset pricing models, Forthcoming in *Review of Financial Studies*.

Gourieroux, Christian, and Joann Jasiak, 2006, Autoregressive gamma processes, *Journal of Forecasting* 25, 129-152.

Gourio, François, 2008, Disasters and recoveries, *American Economic Review Papers and Proceedings* 98, 68–73.

Green, Edward J., 1989, Individual-level randomness in a nonatomic population, working paper, University of Pittsburgh.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, 2014, The nature of countercyclical income risk, *Journal of Political Economy* 122, 621-660.

Hansen, Lars P., John Heaton, and Amir Yaron, 1996, Finite-sample properties of some alternative GMM estimators, *Journal of Business and Economic Statistics* 14, 262-280.

Harvey, Campbell R., and Akhtar Siddique, 2000, Conditional skewness in asset pricing tests, *Journal of Finance* 55, 1263–1295.

Julliard Christian, and Anisha Ghosh, 2012, Can rare events explain the equity premium puzzle? *Review of Financial Studies* 25, 3037-3076.

Krebs, Thomas, 2007, Job displacement risk and the cost of business cycles, *American Economic Review* 97, 664-686.

Kreps, David M., and Evan L. Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46, 185–200.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken, 2010, A skeptical appraisal of asset pricing tests, *Journal of Financial Economics* 96, 175-196.

Mankiw, Gregory N., 1986, The equity premium and the concentration of aggregate shocks, *Journal of Financial Economics* 17, 211-219.

Nakamura, Emi, Jón Steinsson, Robert J. Barro, and José F. Ursúa, 2013, Crises and recoveries in an empirical model of consumption disasters, *American Economic Journal: Macroeconomics* 5, 35-74.

- Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117-131.
- Souleles, Nicholas S., 1999, The response of household consumption to income tax refunds, *American Economic Review* 89, 947-958.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2004, Cyclical dynamics in idiosyncratic labor market risk, *Journal of Political Economy* 112, 695-717.
- Townsend, Robert M., 1994, Risk and insurance in village India, *Econometrica* 62, 539-591.
- Veronesi, Pietro, 2004, The peso problem hypothesis and stock market returns, *Journal of Economic Dynamics and Control* 28, 707–725.
- Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance* 68, 987-1035.
- Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

Table I
Summary Statistics for Household Consumption Growth, Quarterly Data 1982:Q1 to 2009:Q4

The table reports the point estimates of the standard deviation ($\mu_2^{1/2}$), third central moment (μ_3), and fifth percentile of the cross-sectional distribution of quarterly household consumption growth, along with their first-order autocorrelations. Standard errors are in parentheses. The January tranche is the sample of households with first-quarter consumption in January through March; “All tranches” refers to the combined January, February, and March tranches, where the February tranche is the sample of households with first-quarter consumption in February through April and the March tranche is the sample of households with first-quarter consumption in March through May. *AC1* denotes first-order autocorrelation. For the January tranche, the minimum, maximum, and mean number of households in a quarter is 19, 1,310, and 674, respectively. For all tranches, the minimum, maximum, and mean number of households is 64, 3,906, and 2,056, respectively. The time series of the model-implied first, second, and third central moments are obtained by inverting the expressions for the equilibrium price-dividend ratio and risk-free rate to obtain the time series of the state variable and then computing the time series of the cross-sectional moments as affine functions of the state variable (equations (4) to (6)). The time series of the model-implied percentiles are obtained via simulations. Specifically, the time series of the state variable is obtained as above by inverting the expressions for the equilibrium price-dividend ratio and risk-free rate. Then using this aggregate time series, we simulate the consumption growth of 10,000 households of the same length as that of the state variable and obtain the time series of the fifth percentile from the cross-sectional distribution of household consumption growth.

	$\mu_2^{1/2}$	μ_3	Q_{05}	$AC1(\mu_2^{1/2})$	$AC1(\mu_3)$	$AC1(Q_{05})$
Data-Implied Moments						
Jan. tranche	0.379 (0.015)	-0.025 (0.008)	-0.595 (0.026)	0.772 (0.073)	0.124 (0.230)	0.534 (0.583)
All tranches	0.383 (0.015)	-0.024 (0.006)	-0.616 (0.029)	0.897 (0.055)	0.180 (0.146)	0.792 (0.405)
Model-Implied Moments						
	0.219	-0.027	-0.246	0.864	0.876	0.846
Correlation between Data-Implied (All Tranches) and Model-Implied Moments						
	0.490	0.377				

Table II
Correlation between Household Consumption Growth Statistics and Recessions, Quarterly Data
1982:Q1 to 2009:Q4

The table reports the correlation between household consumption growth statistics and recessions. The January tranche is the sample of households with first-quarter consumption in January through March; “All tranches” refers to the combined January, February, and March tranches, where the February tranche is the sample of households with first-quarter consumption in February through April and the March tranche is the sample of households with first-quarter consumption in March through May. I_{rec} is an indicator variable that takes the value of one if there is a NBER-designated recession in at least two of the three months of the quarter. The time series of the model-implied first, second, and third central moments are obtained by inverting the expressions for the equilibrium price-dividend ratio and risk-free rate to obtain the time series of the state variable and then computing the time series of the cross-sectional moments as affine functions of the state variable (equations (4) to (6)). The time series of the model-implied percentiles are obtained via simulations. Specifically, the time series of the state variable is obtained as above by inverting the expressions for the equilibrium price-dividend ratio and risk-free rate. Then, using this aggregate time series, we simulate the consumption growth of 10,000 households of the same length as that of the state variable and obtain the time series of the fifth percentile from the cross-sectional distribution of household consumption growth.

	$corr(\mu_2^{1/2}, I_{rec})$	$corr(\mu_3, I_{rec})$	$corr(Q_{.05}, I_{rec})$
Correlation between Data-Implied Statistics and Recessions			
January tranche	0.139	-0.249	-0.142
All tranches	0.118	-0.217	-0.087
Correlation between Model-Implied Statistics and Recessions			
	0.264	-0.268	-0.233

Table III
Model Fit and Parameter Estimates, Quarterly Data 1982:Q1 to 2009:Q4, January Tranche

The table reports parameter estimates and model fit for the baseline model over the quarterly sample period 1982:Q1 to 2009:Q4. The GMM system consists of 17 moment restrictions (14 aggregate moments and the first three central moments of the cross-sectional distribution of household consumption growth) in 14 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio; and $\mu_2^{1/2}(\Delta c_{CEX})$ and $\mu_3(\Delta c_{CEX})$ denote the volatility and third central moment of the cross-sectional distribution of household consumption growth. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , the EIS ψ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters governing the dynamics of the state variable, ν, ξ , and ρ ; the parameters for household income shocks, $\sigma, \hat{\sigma}$, and $\hat{\omega}$; and the three parameters governing the dynamics of aggregate dividend growth, α_d, β_d , and σ_d . The J -stat is 4.38 and the model is not rejected at the 10% level of significance. The simulated 90%, 95%, and 99% critical values of the J -stat are 6.52, 9.11, and 15.50, respectively. Note that because we use a pre-specified weighting matrix, the J -stat has a nonstandard asymptotic distribution and its critical values are computed via simulation.

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.005	0.005	0.863	0.019	0.087	0.056	3.745	0.411	0.972
	(0.001)	(0.001)	(0.044)	(0.009)	(0.007)	(0.104)	(0.066)	(0.031)	(0.021)
Model	0.007	0.010	0.876	0.020	0.228	-0.061	3.787	0.470	0.876

Consumption and Dividends							
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\mu_2^{1/2}(\Delta c_{CEX})$	$\mu_3(\Delta c_{CEX})$
Data	0.004	0.004	0.005	0.026	0.328	0.379	-0.025
	(0.001)	(0.0004)	(0.003)	(0.004)	(0.088)	(0.016)	(0.008)
Model	0.008	0.004	-0.002	0.025	0.308	0.219	-0.027

Estimates of Preference Parameters		
γ	ψ	δ
1.12	1.00	0.939
(0.396)	(0.526)	(0.150)

Other Parameter Estimates										
μ	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d	$\hat{\omega}$	$\hat{\sigma}$
0.008	0.004	0.040	0.0002	0.876	0.718	0.001	0.020	-40.1	0.110	0.623
(0.001)	(0.001)	(0.324)	(0.006)	(0.125)	(0.027)	(0.102)	(0.613)	(0.001)	(0.085)	(0.161)

Table IV
Model Fit and Parameter Estimates, Quarterly Data 1947:Q1 to 2009:Q4

The table reports parameter estimates and model fit for the baseline model over the quarterly sample period 1947:Q1 to 2009:Q4. The GMM system consists of 13 moment restrictions (only aggregate moments, that is, excluding the first three central moments of the cross-sectional distribution of household consumption growth) in 13 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; and $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , the EIS, ψ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters for the dynamics of the state variable, ν, ξ , and ρ ; the parameters for household income shocks, $\sigma, \hat{\sigma}$, and $\hat{\omega}$; and the three parameters governing the dynamics of aggregate dividend growth, α_d, β_d , and σ_d .

Prices										
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$	
Data	0.003 (0.001)	0.006 (0.001)	0.821 (0.060)	0.017 (0.006)	0.083 (0.005)	0.091 (0.063)	3.470 (0.046)	0.422 (0.029)	0.979 (0.012)	
Model	0.007	0.003	0.960	0.018	0.123	-0.016	3.390	0.437	0.960	
Consumption and Dividends										
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$					
Data	0.005 (0.0004)	0.005 (0.0003)	0.004 (0.002)	0.023 (0.002)	0.369 (0.070)					
Model	0.008	0.004	-0.015	0.037	0.293					
Estimates of Preference Parameter										
γ	ψ	δ								
13.3 (0.001)	1.02 (1.12)	0.988 (0.263)								
Other Parameter Estimates										
μ	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d	$\hat{\omega}$	$\hat{\sigma}$
0.008 (0.001)	0.004 (0.0005)	0.622 (7.03)	0.0007 (0.004)	0.960 (0.133)	0.0000 (6.42)	0.001 (0.168)	0.031 (0.156)	-1.47 (6.23)	0.392 (0.023)	0.039 (0.521)

Table V
Model Fit and Parameter Estimates, Annual Data 1929 to 2009

The table reports parameter estimates and model fit for the baseline model over the annual sample period 1929 to 2009. The GMM system consists of 14 moment restrictions (only aggregate moments, that is, excluding the first three central moments of the cross-sectional distribution of household consumption growth) in 14 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; and $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , the EIS, ψ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters governing the dynamics of the state variable, ν, ξ , and ρ ; the parameters for household income shocks, $\sigma, \hat{\sigma}$, and $\hat{\omega}$; and the three parameters governing the dynamics of aggregate dividend growth, α_d, β_d , and σ_d .

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.006 (0.005)	0.030 (0.005)	0.647 (0.111)	0.070 (0.019)	0.187 (0.019)	-0.084 (0.098)	3.390 (0.081)	0.435 (0.053)	0.864 (0.058)
Model	0.025	0.008	0.837	0.049	0.224	-0.077	3.330	0.399	0.837

Consumption and Dividends					
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$
Data	0.020 (0.002)	0.020 (0.002)	0.012 (0.011)	0.115 (0.018)	0.086 (0.110)
Model	0.018	0.020	0.013	0.054	0.114

Estimates of Preference Parameters		
γ	ψ	δ
5.05 (0.071)	1.10 (0.868)	.983 (0.198)

Other Parameter Estimates										
#	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d	$\hat{\omega}$	$\hat{\sigma}$
0.018 (0.003)	0.020 (0.004)	0.111 (0.538)	0.004 (0.019)	0.837 (0.162)	0.224 (1.54)	0.020 (0.059)	0.050 (0.305)	-2.37 (0.147)	0.448 (0.048)	0.037 (1.19)

Table VI

CRRRA Model Fit and Parameter Estimates, Quarterly Data 1982:Q1 to 2009:Q4, January Tranche

The table reports parameter estimates and model fit for the variant of the baseline model with CRRRA preferences over the quarterly sample period 1982:Q1 to 2009:Q4. The GMM system consists of 17 moment restrictions (14 aggregate moments and the first three central moments of the cross-sectional distribution of household consumption growth) in 13 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio; and $\mu_2^{1/2}(\Delta c_{CEX})$ and $\mu_3(\Delta c_{CEX})$ denote the volatility and third central moment of the cross-sectional distribution of household consumption growth. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters governing the dynamics of the state variable, ν, ξ , and ρ ; the parameters for household income shocks, $\sigma, \hat{\sigma}$, and $\hat{\omega}$; and the parameters governing the dynamics of aggregate dividend growth, α_d, β_d , and σ_d . The J -stat is 14.58 and the model is rejected at the 5% level of significance. The simulated 90%, 95%, and 99% critical values of the J -stat are 6.35, 8.93, and 15.28, respectively. Note that because we use a prespecified weighting matrix, the J -stat has a nonstandard asymptotic distribution and its critical values are computed via simulation.

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.005 (0.001)	0.005 (0.001)	0.863 (0.044)	0.019 (0.009)	0.087 (0.007)	0.056 (0.104)	3.745 (0.066)	0.411 (0.031)	0.972 (0.021)
Model	0.023	0.019	0.912	0.023	0.025	0.125	3.792	0.013	0.912

Consumption and Dividends							
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\mu_2^{1/2}(\Delta c_{CEX})$	$\mu_3(\Delta c_{CEX})$
Data	0.004 (0.001)	0.004 (0.0004)	0.005 (0.003)	0.026 (0.004)	0.328 (0.088)	0.379 (0.016)	-0.025 (0.008)
Model	0.002	0.008	0.001	0.025	0.134	0.059	-0.002

Estimates of Preference Parameters	
γ	δ
2.45 (8.34)	0.943 (0.513)

Other Parameter Estimates										
μ	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d	$\hat{\omega}$	$\hat{\sigma}$
0.002 (0.001)	0.008 (0.001)	0.000 (0.005)	0.096 (5.14)	0.912 (0.132)	0.0007 (0.033)	0.001 (0.012)	0.023 (0.206)	-1.31 (16.3)	0.007 (0.080)	0.672 (1.68)

Table VII
CRRA Model Fit and Parameter Estimates, Annual Data 1929 to 2009

The table reports parameter estimates and model fit for the variant of the baseline model with CRRA preferences over the annual sample period 1929 to 2009. The GMM system consists of 14 moment restrictions (only aggregate moments, that is, excluding the first three central moments of the cross-sectional distribution of household consumption growth) in 13 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; and $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters governing the dynamics of the state variable, ν, ξ , and ρ ; the parameters for household income shocks, $\sigma, \hat{\sigma}$, and $\hat{\omega}$; and the parameters governing the dynamics of aggregate dividend growth, α_d, β_d , and σ_d . The J -stat is 26.02 and the model is rejected at the 1% level of significance. The simulated 90%, 95%, and 99% critical values of the J -stat are 10.11, 14.34, and 24.78, respectively. Note that because we use a pre-specified weighting matrix, the J -stat has a nonstandard asymptotic distribution and its critical values are computed via simulation.

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.006	0.030	0.647	0.070	0.187	-0.084	3.390	0.435	0.864
	(0.005)	(0.005)	(0.111)	(0.019)	(0.019)	(0.098)	(0.081)	(0.053)	(0.058)
Model	0.042	0.029	0.787	0.036	0.114	0.117	3.397	0.069	0.787

Consumption and Dividends					
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$
Data	0.020	0.020	0.012	0.115	0.086
	(0.002)	(0.002)	(0.011)	(0.018)	(0.110)
Model	0.011	0.023	0.003	0.112	0.139

Estimates of Preference Parameters	
γ	δ
14.7	0.987
(0.242)	(0.002)

Other Parameter Estimates										
μ	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d	$\hat{\omega}$	$\hat{\sigma}$
.011	.023	.001	.167	.787	.00001	.005	.110	-1.75	.951	.026
(0.003)	(0.004)	(0.056)	(3.39)	(0.171)	(0.016)	(0.050)	(0.164)	(8.04)	(0.000)	(0.012)

Table VIII
Lognormal Model Fit and Parameter Estimates,
Quarterly Data 1982:Q1 to 2009:Q4, January Tranche

The table reports parameter estimates and model fit for the variant of the baseline model with lognormal shocks over the quarterly sample period 1982:Q1 to 2009:Q4. The GMM system consists of 16 moment restrictions (14 aggregate moments and the first two central moments of the cross-sectional distribution of household consumption growth) in 12 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio; and $\mu_2^{1/2}(\Delta c_{CEX})$ and $\mu_3(\Delta c_{CEX})$ denote the volatility and third central moment of the cross-sectional distribution of household consumption growth. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , the EIS, ψ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters for the dynamics of the state variable, ν_σ , ξ_σ , and ρ_σ ; the parameter for household income shocks, $\hat{\sigma}$; and the parameters governing the dynamics of aggregate dividend growth, α_d , β_d , and σ_d . The J -stat is 4.63 and the model is not rejected at the 10% level of significance. The simulated 90%, 95%, and 99% critical values of the J -stat are 6.30, 8.85, and 15.09, respectively. Note that because we use a pre-specified weighting matrix, the J -stat has a nonstandard asymptotic distribution and its critical values are computed via simulation.

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.005 (0.001)	0.005 (0.001)	0.863 (0.044)	0.019 (0.009)	0.087 (0.007)	0.056 (0.104)	3.745 (0.066)	0.411 (0.031)	0.972 (0.021)
Model	0.020	0.006	0.922	0.021	0.163	-0.032	3.787	0.419	0.922

Consumption and Dividends							
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\mu_2^{1/2}(\Delta c_{CEX})$	$\mu_3(\Delta c_{CEX})$
Data	0.004 (0.001)	0.004 (0.000)	0.005 (0.003)	0.026 (0.004)	0.328 (0.088)	0.379 (0.016)	-0.025 (0.008)
Model	0.007	0.004	-0.001	0.037	0.314	0.297	0.0

Estimates of Preference Parameters		
γ	ψ	δ
1.03 (14.5)	1.00 (0.501)	0.901 (1.13)

Other Parameter Estimates								
μ	σ_a	ν_σ	ξ_σ	ρ_σ	α_d	σ_d	β_d	$\hat{\sigma}$
0.007 (0.001)	0.004 (0.001)	0.008 (0.104)	0.003 (0.012)	0.922 (0.125)	0.001 (0.024)	0.030 (0.160)	-7.25 (0.388)	0.297 (0.024)

Table IX
Model Fit and Parameter Estimates, Quarterly Data 1982:Q1 to 2009:Q4,
January Tranche: Only Business Cycle Shock

The table reports parameter estimates and model fit for the baseline model over the quarterly sample period 1982:Q1 to 2009:Q4. The GMM system consists of 17 moment restrictions (14 aggregate moments and the first three central moments of the cross-sectional distribution of household consumption growth) in 12 parameters. $E[r_f]$, $\sigma(r_f)$, and $AC1(r_f)$ are the mean, standard deviation, and first-order autocorrelation of the risk-free rate; $E[r_m]$, $\sigma(r_m)$, and $AC1(r_m)$ are the mean, standard deviation, and first-order autocorrelation of the market return; $E[p/d]$, $\sigma(p/d)$, and $AC1(p/d)$ are the mean, standard deviation, and first-order autocorrelation of the price-dividend ratio; and $\mu_2^{1/2}(\Delta c_{CEX})$ and $\mu_3(\Delta c_{CEX})$ denote the volatility and third central moment of the cross-sectional distribution of household consumption growth. Δc is aggregate consumption growth and Δd is dividend growth. The preference parameters are the RRA coefficient, γ , the EIS, ψ , and the subjective discount factor, δ . The other parameters are: the mean, μ , and volatility, σ_a , of aggregate consumption growth; the parameters governing the dynamics of the state variable, ν , ξ , and ρ ; the parameter for household income shocks, σ ; and the parameters governing the dynamics of aggregate dividend growth, α_d , β_d , and σ_d . The J -stat is 5.84 and the model is not rejected at the 10% level of significance. The simulated 90%, 95%, and 99% critical values of the J -stat are 6.20, 8.74, and 14.99, respectively. Note that because we use a pre-specified weighting matrix, the J -stat has a nonstandard asymptotic distribution and its critical values are computed via simulation.

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	0.005 (0.001)	0.005 (0.001)	0.863 (0.044)	0.019 (0.009)	0.087 (0.007)	0.056 (0.104)	3.745 (0.066)	0.411 (0.031)	0.972 (0.021)
Model	-0.008	0.023	0.992	0.012	0.055	0.017	3.719	0.386	0.992

Consumption and Dividends							
	$E[\Delta c]$	$\sigma(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\mu_2^{V2}(\Delta c_{CXX})$	$\mu_3(\Delta c_{CXX})$
Data	0.004 (0.001)	0.004 (0.000)	0.005 (0.003)	0.026 (0.004)	0.328 (0.088)	0.379 (0.016)	-0.025 (0.008)
Model	0.008	0.005	-0.012	0.032	0.277	0.130	-0.022

Estimates of Preference Parameters		
γ	ψ	δ
1.05 (0.278)	1.00 (0.046)	0.989 (0.068)

Other Parameter Estimates								
#	σ_a	ν	ξ	ρ	σ	α_d	σ_d	β_d
0.008 (0.001)	0.005 (0.390)	1.355 (0.331)	0.000 (0.000)	0.992 (0.129)	0.999 (0.641)	0.008 (0.113)	0.027 (0.182)	-58.13 (0.279)

Table X
Rolling Fama-MacBeth Regressions, Quarterly Data 1982:Q1 to 2009:Q4, January Tranche

The table reports Fama-Macbeth (1973) cross-sectional regression results using as test assets the 25 Fama-French (1993) portfolios (Panel A), the 30 industry-sorted portfolios (Panel B), and the combined set of 25 Fama-French and 30 industry-sorted portfolios (Panel C). The data are quarterly over 1982:Q1 to 2009:Q4. The adjusted R squared, \overline{R}^2 , are reported. The standard errors of \hat{a} and $\hat{\lambda}$ are calculated from the time series of the cross-sectional intercepts and slope coefficients. The factor loadings are estimated at each period t , starting with the midpoint of the sample, using all the returns up to period t .

\overline{R}^2	$\hat{\alpha}$	$\hat{\lambda}_{skew}$	$\hat{\lambda}_{std}$	$\hat{\lambda}_{MKT}$	$\hat{\lambda}_{SMB}$	$\hat{\lambda}_{HML}$
Panel A: 25 FF Portfolios						
13.6%	0.01 (0.01)	0.73 (0.61)				
40.0%	0.01 (0.01)		-0.09 (0.07)			
37.4%	0.01 (0.01)	0.01 (0.70)	-0.10 (0.07)			
59.5%	0.04 (0.02)			-0.03 (0.03)	0.01 (0.01)	0.01 (0.01)
Panel B: 30 Industry Portfolios						
14.0%	0.01 (0.01)	0.34 (0.38)				
-6.9%	0.01 (0.01)		0.00 (0.07)			
10.4%	0.02 (0.01)	0.74 (0.38)	0.09 (0.07)			
-22.8%	0.02 (0.02)			-0.01 (0.02)	-0.01 (0.01)	0.005 (0.01)
Panel C: 25 FF and 30 Industry Portfolios						
14.9%	0.01 (0.01)	0.37 (0.37)				
5.0%	0.02 (0.01)		-0.03 (0.06)			
9.4%	0.02 (0.01)	0.54 (0.39)	0.01 (0.06)			
-7.5%	0.03 (0.02)			-0.01 (0.02)	0.004 (0.01)	0.01 (0.01)

Table XI
Fixed Fama-MacBeth Regressions,
Quarterly Data 1982:Q1 to 2009:Q4, January Tranche

The table reports Fama-Macbeth (1973) cross-sectional regression results using as test assets the 25 Fama-French (1993) portfolios (Panel A), the 30 industry-sorted portfolios (Panel B), and the combined set of 25 Fama-French and 30 industry-sorted portfolios (Panel C). The data are quarterly over 1982:Q1 to 2009:Q4. The adjusted R squared, \overline{R}^2 , are reported. The standard errors of \hat{a} and $\hat{\lambda}$ are calculated from the time series of the cross-sectional intercepts and slope coefficients. The factor loadings are estimated on the first half of the sample.

\overline{R}^2	$\hat{\alpha}$	$\hat{\lambda}_{skew}$	$\hat{\lambda}_{std}$	$\hat{\lambda}_{MKT}$	$\hat{\lambda}_{SMB}$	$\hat{\lambda}_{HML}$
Panel A: 25 FF Portfolios						
21.5%	0.01 (0.01)	0.92 (0.65)				
42.8%	0.01 (0.01)		-0.11 (0.06)			
41.6%	0.01 (0.01)	0.16 (0.90)	-0.11 (0.06)			
53.6%	0.04 (0.03)			-0.03 (0.03)	0.01 (0.01)	0.01 (0.01)
Panel B: 30 Industry Portfolios						
7.5%	0.01 (0.01)	0.33 (0.39)				
7.9%	0.02 (0.01)		0.06 (0.06)			
39.0%	0.02 (0.01)	0.57 (0.42)	0.10 (0.06)			
28.3%	0.07 (0.03)			-0.05 (0.03)	0.002 (0.01)	0.001 (0.01)
Panel C: 25 FF and 30 Industry Portfolios						
9.8%	0.01 (0.01)	0.40 (0.37)				
-2.0%	0.02 (0.01)		0.00 (0.05)			
13.2%	0.02 (0.01)	0.53 (0.41)	0.02 (0.05)			
30.2%	0.07 (0.03)			-0.06 (0.02)	0.004 (0.01)	0.01 (0.01)

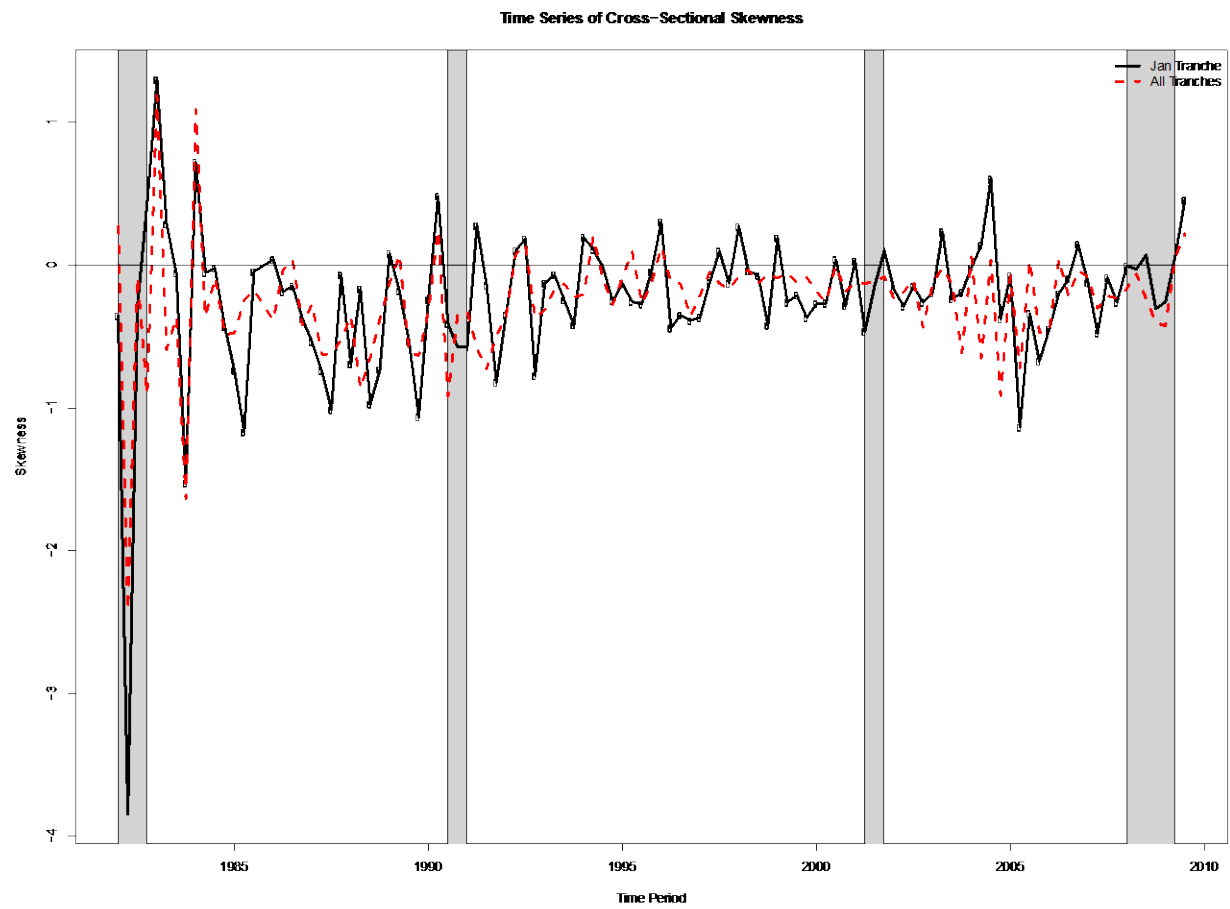


Figure 1. Time series of the cross-sectional skewness, quarterly data 1982:Q1 to 2009:Q4.

¹ In a different context, in particular, relying on a non-parametric relative entropy minimizing approach to filter the most likely SDF, Ghosh, Julliard, and Taylor (2014) highlight the importance of higher moments of the SDF, especially skewness, in pricing assets. In particular, they show that about a quarter of the overall entropy of the most likely SDF is generated by its third- and higher-order moments, with the third central moment alone accounting for about 18% of the entropy.

² Related references include Backus, Chernov, and Martin (2011), Barro and Ursua (2008), Constantinides (2008), Gabaix (2012), Gourio (2008), Harvey and Siddique (2000), Julliard and Ghosh (2012), Nakamura et al. (2013), Veronesi (2004), and Wachter (2013).

³ The probability distribution of the random variable $j_{i,s}^{1/2} \sigma \eta_{i,s}$ is known as a Poisson mixture of normals. This distribution is tractable because it is normal, conditional on $j_{i,s}$.

⁴ The argument is due to Green (1989) and is elaborated in Appendix A.

⁵ Recursive preferences were introduced by Kreps and Porteus (1978) and adapted in the form used here by Epstein and Zin (1989) and Weil (1990).

⁶ Essentially, we build into the model the assumption that the consumption growth of all households in a given period is independent of each household's consumption level. A richer model would allow the consumption growth of each household in a given period to depend on the household's consumption level, consistent with the empirical findings of Guvenen, Ozkan, and Song (2014). Guvenen, Ozkan, and Song (2014) analyze the confidential earnings histories of millions of individuals over the period 1978 to 2010 and find that the earning power of the lowest income workers and the top 1% of income workers declines the most in recessions, compared to other workers.

⁷ The reader may wonder why the model-implied third central moment is always negative. While the third central moment of $j_{i,s}^{1/2} \sigma \eta_{i,s}$ is zero,

$$\mu_3 \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} \right) \mid \omega_{t+1} \right] = E \left[E \left[\left(j_{i,t+1}^{1/2} \sigma \eta_{i,t+1} \right)^3 \mid j_{i,t+1} \right] \mid \omega_{t+1} \right] = 0,$$
 the third central moment of $-j_{i,s} \sigma^2 / 2$ is negative and this imparts a negative third central moment to the random variable $j_{i,s}^{1/2} \sigma \eta_{i,s} - j_{i,s} \sigma^2 / 2$.

⁸ Equations (4) to (6) show that the mean, variance, and third central moment of the cross-sectional distribution of relative household consumption growth are affine in the single state variable, household consumption risk, ω_t . Therefore, an increase in household consumption risk simultaneously results in a decrease in the cross-sectional mean and an increase in both the cross-sectional variance and the absolute value of the (negative) third central moment.

⁹ The ARG (1) process is the exact discrete-time equivalent of the square-root (CIR) process, and is defined as follows (see, for example, Gourioux and Jasiak (2006)): $(x_{t+1} / \xi) \mid y_{t+1} \sim \Gamma(\nu + y_{t+1}), \nu > 0, y_{t+1} \mid x_t \sim P(\rho x_t / \xi), \rho, \xi > 0$, where Γ denotes a gamma distribution, P denotes a Poisson distribution, ξ is a scale parameter, ν is the degree of freedom, ρ is the correlation parameter, and y_{t+1} is the mixing variable. The conditional probability density function of an ARG (1) process, $f(x_{t+1} \mid x_t; \nu, \xi, \rho)$, is a mixture of gamma densities with

Poisson weights. Therefore, the ARG (1) process is strictly positive. Moreover, it also admits the autoregressive representation in equation (7).

¹⁰ We draw a distinction between the stock market and the “market,” which we define earlier as the sum total of all assets in the economy. The log dividend growth of the stock market is Δd_{t+1} .

¹¹ Our description and filters of the household consumption data closely follow Brav, Constantinides, and Geczy (2002).

¹² If we were to exclude the training quarter in classifying a household as in the panel, then each household would stay in the panel for four quarters and new households would replace one-fourth of the participating households each quarter.

¹³ The constant rotation of the panel makes it impossible to test hypotheses regarding a specific household’s behavior over time for more than four quarters. A longer time series of individual household consumption is available from the PSID database, but only for food consumption.

¹⁴ See Attanasio and Weber (1995) and Souleles (1999) for further details regarding the database.

¹⁵ A relatively mild minimum asset criterion of \$2,000 for a household to be included in the sample eliminates about 80% of households and eliminates all households in some quarters. Stricter filters further eliminate households to the point that statistics with a small number of households become unreliable. In the interest of having a large sample, we do not impose a minimum asset filter in the reported tests.

¹⁶ See Table 4 of Blundell, Pistaferri, and Preston (2008).

¹⁷ This equality holds if the measurement error has a stationary cross-sectional distribution or if the higher-order (>2) moments of its cross-sectional distribution are zero for all time periods t .

¹⁸ The pre-specified weighting matrix has two advantages over the efficient weighting matrix. First, it has superior small-sample properties (see, for example, Ahn and Gadarowski (1999), Ferson and Foerster (1994), and Hansen, Heaton, and Yaron (1996)). Second, the moment restrictions included in the GMM have different orders of magnitude, with the mean of the price-dividend ratio being a couple of orders of magnitude larger than the means of the market return and risk-free rate. Therefore, placing larger weights on the latter two moments allows the GMM procedure to put equal emphasis in matching all these moments. We repeated our estimation using the efficient weighting matrix and obtained similar results.

¹⁹ Similar results are obtained for the February and March tranches.

²⁰ The model implies that the risk-free rate and price-dividend ratio are affine functions of the state variable. We use the point estimates of the parameters and extract the current value of the state variable from the observed risk-free rate and price-dividend ratio by minimizing the least-squares criterion function. Given the current value of the state variable, we calculate the model-implied cross-sectional moments using equations (4) to (6).

²¹ See Lewellen, Nagel, and Shanken (2010).

²² Note that equation (D3) implies that A_1 is the solution of a quadratic. We verified, via simulations, that the economically meaningful root is the smaller of the two.

²³ Note that equation (D10) implies that B_1 is the solution of a quadratic equation. We verified, via simulations, that the economically meaningful root is the smaller of the two.