NBER WORKING PAPER SERIES

CONSUMPTION RISK AND CROSS-SECTIONAL RETURNS

Jonathan A. Parker
Christian Julliard

Working Paper 9538
http://www.nber.org/papers/w9538

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2003

For helpful comments and discussions, we thank Yacine Ait-Sahalia, Albina Danilova, Pierre-Oliver Gourinchas, Sydney Ludvigson, Monika Piazzesi, Christopher Sims and seminar participants at Princeton University. For financial suport, Parker thanks the National Science Foundation grant SES-0096076, and the Sloan Foundation. Contact information and revised versions available at:www.princeton.edu/~jparker/.The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Consumption Risk and Cross-Sectional Returns
Jonathan A. Parker and Christian Julliard
NBER Working Paper No. 9538
March 2003
JEL No. G12, G11, E21

ABSTRACT

This paper evaluates the central insight of the Consumption Capital Asset Pricing Model (C-CAPM) that an asset's expected return is determined by its equilibrium risk to consumption. Rather that measure the risk of a portfolio by the contemporaneous covariance of its return and consumption growth -- as done in the previous literature on the C-CAPM and the pattern of cross-sectional returns -- we measure the risk of a portfolio by its ultimate consumption risk defined as the covariance of its return and consumption growth over the quarter of the return and many following quarters. While contemporaneous consumption risk has little predictive power for explaining the pattern of average returns across the Fama and French (25) portfolios, ultimate consumption risk is highly statistically significant in explaining average returns and explains a large fraction of the variation in average returns. Aditionally, estimates of the average risk-free real rate of interest and the coefficient of relative risk aversion of the representative household based on ultimate consumption risk are more reasonable than those obtained using contemporaneous consumption risk.

Jonathan A. Parker
Department of Economics
Princeton University
Princeton, NJ 08544
and NBER
jparker@princeton.edu

Christian Julliard
Department of Economics
Princeton University
Princeton, NJ 08544
julliard@princeton.edu

1. Introduction

The natural economic explanation for the large predictable differences in returns observed across assets in the U.S. stock market is differences in risk. According to canonical economic theory, an asset's risk is its covariance with consumption growth. But observed differences in the covariance of returns and consumption growth across portfolios do not explain the observed differences in expected returns across portfolios.[1] The asset pricing literature has largely concluded that the textbook consumption-based capital asset pricing model (C-CAPM) does not explain observed differences in returns across portfolios of stocks, and instead differences in return arise from time-variation in effective risk aversion or quite different models of economic behavior.

In this paper, we study the Fama and French size and book-to-market portfolios and re-evaluate the central insight of the Consumption Capital Asset Pricing model (C-CAPM) that an asset's expected return is determined by its equilibrium risk to consumption. Rather that measure the risk of a portfolio by the contemporaneous covariance of its return and consumption growth – as done in the previous literature on the C-CAPM and the pattern of cross-sectional returns – we follow Parker (2001) and measure the risk of a portfolio by its ultimate risk to consumption, defined as the covariance of its return and consumption growth over the quarter of the return and many following quarters.

Measuring the risk of equity as the ultimate impact of a return on consumption has several appealing features. First, this approach maintains the assumption that the primary determinant of utility is the level of flow consumption. This assumption is intuitive and has proved useful and successful in many other subfields of economics. Second, this approach is consistent with the textbook model of portfolio choice in that, if the textbook model were true, the ultimate risk would correctly measure the risk of different portfolios. Finally and most importantly, the ultimate risk is a better measure of the true risk of the stock market under a wide class of extant models used in the study of household consumption and saving. If consumption responds with a lag to changes in wealth, then the contemporaneous covariance of consumption and wealth understates the risk of equity, and the ultimate risk provides the correct measure. This slow adjustment is a well-documented feature of consumption data: consumption displays excess smoothness in response to wealth shocks.[2] Existing theoretical explanations for slow adjustment of consumption include direct costs of adjusting consumption, nonseparability of the marginal utility of consumption from factors such as labor supply or housing stock, which themselves are constrained to adjust slowly, constraints on borrowing or changes in risk that hinder consumption smoothing, and constraints on information flow or calculation so that household be-

[1] See Mankiw and Shapiro (1986), Breeden, Gibbons, and Litzenberger (1989), Campbell (1996), Cochrane (1996) and Lettau and Ludvigson (2001b).

[2] The slow adjustment of consumption has a long history starting with Flavin (1981) and Hall and Mishkin (1982). Even the seminal paper, Hall (1978), which first notes that consumption should not be slow to adjust rejects the random walk of consumption using data on market returns.

havior is "near-rational." The common feature of these models is that consumption responds slowly to a market innovation, so that only after some time is the impact of a change in wealth completely observed in the movement of consumption.

Thus, the ultimate risk provides a robust measure of the risk of a stock in that it remains to some extent agnostic about the particular optimization problem faced by households. This robustness allows us to evaluate the economic insight that consumption risk should determine returns even though the true or complete model of household saving and portfolio choice has to date escaped discovery.

We find three main pieces of evidence that consumption risk largely explains the cross-sectional pattern of returns.

First, while, the covariance of each portfolio and contemporaneous consumption growth has almost no predictive power for explaining the pattern of average returns across portfolios, the ultimate risk to consumption explains half or more of the variation in returns across portfolios of stocks. The explanatory power of this single factor – ultimate consumption risk –is similar to that of the three factors models of Fama and French (1993) and Lettau and Ludvigson (2001). Second, and related, the ultimate risk to consumption is statistically significant while the contemporaneous risk is not. We also show that firm characteristics do not drive out the statistical significance of consumption risk. These findings are consistent with the slow adjustment of consumption to a change in wealth and inconsistent with the textbook model in which consumption adjusts instantaneously to a return.

Third, the baseline model implies a risk-free real interest rate around five percent, a point estimate that is not statistically different from the average real rate on 3-month Treasury bills in the sample. While some alternative specifications yield estimates of the risk-free rate that are lower and closer to the sample average, even this estimate is a significant improvement both over estimates based on contemporaneous consumption risk and over estimates based on the macroeconomic factor model of Lettau and Ludvigson (2001b). Finally, we use the model to calculate the risk aversion of the representative agent from the estimated relationship between risk and return. Point estimates of risk aversion lie around 12, again much lower than average estimates implied by the models of Lettau and Ludvigson (2001b), Campbell and Cochrane (1999) or contemporaneous consumption risk, which imply levels of risk aversion of 40 and larger. Using the ultimate risk to consumption, a relatively-small 95 percent confidence interval contains reasonable estimates of risk aversion.

Our main results are most closely related to Brainard, Nelson, and Shapiro (1991) who show that the longer the horizon of the investor, the better the C-CAPM performs relative to the CAPM and to Bansal, Dittmar, and Lundblad (2001) who show that the cointegrating relationship between consumption and dividends explains a large share of the variation in average returns. We work directly with returns rather than long-run movements in dividends and map estimates back to underlying structural parameters which gives us confirmation of the importance of consumption risk. We are also closely related to the literature on the stochastic properties of aggregate consumption following aggregate market returns (Daniel and Marshall (1997), Kandel and Stambaugh (1990), Parker (1999b), Ludvigson and Steindel (1999), Parker (2001), Dynan and Maki (2001), Gabaix and Laibson (2001), and Piazzesi (2001)).

The balance of the paper is organized as follows. The next section sets the stage by deriving the textbook C-CAPM from the portfolio choice of a representative agent. Section 3 presents the ultimate risk to consumption and discusses the range of models that would imply it. The fourth section describes our data and the fifth our econometric methodology. The sixth section contains our main results on fit, significance, risk-free rate, and implied risk aversion for the ultimate risk to consumption. The seventh section of the paper compares our one-factor model to the fit and the performance of the models of Fama and French (1993) and Lettau and Ludvigson (2001b). The eighth section presents results using the ultimate risk to consumption in an economy with time variation in the risk-free rate of interest. A final section concludes.

2. The average cross-sectional pattern of returns and the C-CAPM

Following the literature, we use the equilibrium condition governing optimal portfolio choice to derive a beta representation of average returns across assets as a function of their consumption risk. This derivation of the beta representation of the C-CAPM lays the groundwork for the derivation of our beta representation based on the ultimate risk to consumption.

The C-CAPM first developed by Rubinstein (1976) and Breeden (1979) assumes that the representative household maximizes the expected present discounted value of utility flows from consumption by allocating wealth to consumption and different investment opportunities. At the optimal allocation, a small extra investment in a portfolio of stocks yields an expected marginal increase in utility at $t+1$ that exactly offsets the expected marginal decrease in utility from the small amount less invested in the risk-free asset:

$$E_t \left[u' \left(C_{t+1}\right) R_{i,t+1}\right] - E_t \left[u' \left(C_{t+1}\right)\right] R^f_{t,t+1} = 0. \tag{2.1}$$

where $u(\cdot)$ is the period utility function, C is consumption, $R^f_{t,t+1}$ is the gross risk-free real interest rate in the economy between t and $t+1$, and $R_{i,t+1}$ is the gross return on portfolio i of stocks, unknown at t, known at $t+1$.

Equation (2.1) can be written as a model of average cross-sectional returns by manipulating it to a beta representation or factor model. First, we assume that the risk-free rate in the economy is constant. This assumption allows us to relate our results to previous work and keep the algebra simple, but more importantly is a reasonable approximation since, relative to the portfolios of stocks that we consider, the risk-free rate moves little. Section 8 confirms that relaxing this assumption in the textbook model does not affect the substance of our results. Second, to make the model empirically tractable, divide $u'(C_{t+1})$ by $u'(C_t)$ to move from the nonstationary $u'(C_{t+1})$ to the stationary $u'(C_{t+1})/u'(C_t)$. Third, take the unconditional expectation, re-write the expectation of the product in terms of covariances, and re-organize to yield

$$E\left[R^i_{t+1}\right] = \alpha_0 + \beta_{i,0}\lambda_0 \tag{2.2}$$

where

$$\alpha_0 = R^f,\ \beta_{i,0} = \frac{Cov\left[\frac{u'(C_{t+1})}{u'(C_t)}, R_{i,t+1}\right]}{Var\left[\frac{u'(C_{t+1})}{u'(C_t)}\right]},\ \lambda_0 = -\frac{Var\left[\frac{u'(C_{t+1})}{u'(C_t)}\right]}{E\left[\frac{u'(C_{t+1})}{u'(C_t)}\right]}$$

This equation summarizes the implications of the canonical C-CAPM for cross-sectional returns. The expected return on a portfolio is the risk-free rate of return plus the scaled consumption risk of the portfolio, $\beta_{i,0}\lambda_0$. Since $\lambda_0 < 0$, portfolios that pay off poorly when marginal utility is high (so consumption is scarce) have a high expected return.

To proceed to estimation and testing of this model of cross-sectional returns, we log-linearize the ratio of marginal utilities assuming constant relative risk aversion, γ, as

$$\frac{u'(C_{t+1})}{u'(C_t)} \approx 1 - \gamma\Delta\ln C_{t+1}$$

and treat consumption growth as the stochastic discount factor that prices returns.[3] Including the preference parameter, γ, in the coefficient λ, the model for returns is still in the form of equation (2.2) but with coefficients

$$\alpha_0 = E\left[R^f\right],\ \beta_{i,0} = \frac{Cov\left[\Delta\ln C_{t+1}, R_{i,t+1}\right]}{Var\left[\Delta\ln C_{t+1}\right]},\ \lambda_0 = \frac{\gamma Var\left[\Delta\ln C_{t+1}\right]}{E\left[1-\gamma\Delta\ln C_{t+1}\right]} \tag{2.3}$$

Linearization is not necessary, but it allows us to use the Fama and MacBeth (1973) two-step procedure to evaluate the model and estimate γ, which in turn has the advantage of making our results easy to understand and easy to relate to previous results.

We estimate and test equation (2.2). Further, equation (2.3) provides an external test of the structure embodied in the model. The estimated α should equal the risk-free rate of return and the estimated λ_0 and moments of consumption growth should imply a reasonable level of the risk aversion for the representative investor, according to

$$\gamma = \frac{\lambda_0}{E\left[\Delta\ln C_{t+1}\right]\lambda_0 + Var\left[\Delta\ln C_{t+1}\right]}. \tag{2.4}$$

[3] This stochastic discount factor only prices returns relative to one another and not goods over time because we derive it from portfolio choice rather than the intertemporal allocation of consumption. That is, our ratio of marginal utilities omits the household's degree of impatience (the household's discount factor).

3. The ultimate risk to consumption

Equations (2.1), (2.3), and (2.4) evaluate the risk of a portfolio based solely on its covariance with contemporaneous consumption growth. They maintain the assumption that the intertemporal allocation of consumption is optimal from the perspective of the textbook model of consumption smoothing, so that any change in marginal utility is reflected instantly and completely in consumption. This is at odds with the literature on the intertemporal allocation of consumption, which is in wide agreement that the simple, textbook model of a representative consumer is false.[4] In particular, the empirical literature studying consumption and saving behavior suggests that the following assumptions are at least questionable and at worst quite misleading: 1) utility is additively separable from factors that adjust slowly and covary with returns, such as leisure; 2) uninsurable idiosyncractic risk and borrowing constraints are not important; 3) consumption can be instantaneously adjusted or if there are adjustment costs on some items, such as housing, then these categories are additively separable from the utility of consumption; 4) aggregate consumption data accurately measure movements in flow consumption; 5) households perfectly optimize without informational or calculation constraints.[5]

To allow for the slow response of consumption to market returns, this paper evaluates the risk/return trade-off among portfolios of stocks by focussing on the ultimate risk to consumption. Rather that measure the risk to consumption from the contemporaneous comovement of consumption and returns, the risk to consumption is measured by the response of consumption to a return over a longer horizon, as given by

$$Cov\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), R_{i,t+1}\right] \tag{3.1}$$

or in beta representation

$$E\left[R^i\right] = \alpha_S + \beta_{i,S}\lambda_S \tag{3.2}$$

where

$$\alpha_S = R^f,\ \beta_{i,S} = \frac{Cov\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right), R_{i,t+1}\right]}{Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]},\ \lambda_S = \frac{\gamma Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}{E\left[\left(1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right)\right]} \tag{3.3}$$

Thus the stochastic discount factor we consider is one minus the long-horizon consumption growth times the risk aversion of the representative agent. Consumption

[4] This literature is not in agreement about the explanation for this failure, but this literature has the problem of too many models fitting time series data on consumption and risk-free returns rather than no model fitting this behavior.

[5] See Attanasio and Weber (1995), Basu and Kimball (2000), Zeldes (1989), Caballero (1990), Carroll (1997), Gourinchas and Parker (2002), Grossman and Laroque (1990), Ogaki and Reinhart (1998), Attanasio and Weber (1995), Baxter and Jermann (1999), Attanasio and Weber (1993), Wilcox (1992), Parker (1999a), Souleles (forthcoming).

risk is measured by the covariance of the return at $t+1$ and the change in consumption from t to $t+1+S$, where S is the horizon over which the consumption response is studied. The implied relative risk aversion of the representative agent is

$$\gamma_S = \frac{\lambda_S}{E\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]\lambda_S + Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]} \tag{3.4}$$

Why evaluate risk using ultimate risk, as in equations (3.3) and (3.4)? If households choose their portfolio at time t, and the impact of this choice and the realized return on stocks takes time to appear in observed consumption data, then this measure provides a better measure of the risks of stocks than does the contemporaneous measure. To be more concrete, the remainder of this subsection argues that the ultimate consumption risk measures consumption risk *a*) if the textbook model is true *b*) if the marginal utility of consumption is shifted by a stationary variable that covaries with returns, *c*) if constraints on information flow slow consumption movements, and *d*) if consumption data is mismeasured due to the use of lagged data.

First consider a textbook model of consumption smoothing. Households seek to smooth consumption over time, which is captured by the consumption Euler equation between any two periods

$$E_\tau\left[\frac{\delta R^f u'(C_{\tau+1})}{u'(C_\tau)}\right] = 1$$

where δ is the factor by which households discount the future. These relationships imply

$$u'(C_{t+1}) = \left(\delta R^f\right)^S u'(C_{t+1+S}) - \sum_{\tau=t+2}^{t+1+S}\left(\delta R^f\right)^{\tau-t-1}\varepsilon_\tau$$

where $\varepsilon_\tau = u'(C_\tau) - E_{\tau-1}[u'(C_\tau)]$, the expectation error in τ. Substituting into equation (2.1) and noting that the expectation errors are mean zero and uncorrelated with information known at time $t+1$ yields

$$E_t[u'(C_{t+1+S})R_{i,t+1}] - E_t[u'(C_{t+1+S})]R^f = 0 \tag{3.5}$$

Following the same derivation as in section (2) yields the structural beta representation given by equations (3.2) to (3.4). Thus, if the textbook model were true, our analysis would be as valid as the analysis using contemporaneous risk.

Second, consider the class of models in which marginal utility adjusts at the time of the return, but in which a stationary confounding variable implies that the contemporaneous change in consumption understates this change. For this class of models, the ultimate risk to consumption explains the pattern of expected returns, but our estimates of risk aversion are biased. More concretely, suppose that the marginal utility of the representative agent is shifted by a stationary variable, $\mu_t > 0$, so that marginal utility is given by $\mu_t u'(C_t)$, and normalize so that $\lim_{S\to\infty} E_{t+1}\left[\mu_{t+1+S}\right] = 1$.

Assume that the η_t^u,μ_t is a jointly covariance stationary process, where $\eta_t^u = u'(C_t) - E_{t-1}[u'(C_t)]$.[6] μ_t can capture the factors that shift the marginal utility of consumption in many extant models. For example, μ can represent a) the share of hours devoted to market work, b) the relative productivity of home production, c) the stock of durable goods relative to flow consumption.[7]

In such a model, arbitrage across assets implies

$$E_t\left[\mu_{t+1}u'(C_{t+1})R_{i,t+1}\right] = E_t\left[\mu_{t+1}u'(C_{t+1})\right]R^f$$

and consumption smoothing implies

$$\mu_{t+1}u'(C_{t+1}) = \left(\delta R^f\right)^S \mu_{t+1+S}u'(C_{t+1+S}) - \sum_{\tau=t+2}^{t+1+S}\left(\delta R^f\right)^{\tau-t-1}\varepsilon_\tau.$$

Combining gives

$$\begin{aligned} E_t\left[\mu_{t+1+S}u'(C_{t+1+S})\right]R^f &= E_t\left[\mu_{t+1+S}u'(C_{t+1+S})R_{i,t+1}\right] \\ &= E_t\left[E_{t+1}\left[\mu_{t+1+S}u'(C_{t+1+S})\right]R_{i,t+1}\right]. \end{aligned}$$

Now, defining $\psi = \lim_{S\to\infty} Cov_{t+1}\left[\sum_{\tau=S/2+1}^{S}\eta^u_{t+1+\tau}, \mu_{t+1+S}\right]$, we have

$$\begin{aligned} E_{t+1}\left[\mu_{t+1+S}u'(C_{t+1+S})\right] &= Cov_{t+1}\left[\mu_{t+1+S}, u'(C_{t+1+S})\right] + E_{t+1}\left[\mu_{t+1+S}\right]E_{t+1}\left[u'(C_{t+1+S})\right] \\ &\underset{S\to\infty}{\longrightarrow} \psi + E_{t+1}\left[u'(C_{t+1+S})\right] \end{aligned}$$

since

$$\begin{aligned} Cov_{t+1}\left[\mu_{t+1+S}, u'(C_{t+1+S})\right] &= Cov_{t+1}\left[\sum_{\tau=1}^{S/2}\eta^u_{t+1+\tau}, \mu_{t+1+S}\right] + Cov_{t+1}\left[\sum_{\tau=S/2+1}^{S}\eta^u_{t+1+\tau}, \mu_{t+1+S}\right] \\ &\underset{S\to\infty}{\longrightarrow} 0 + \psi. \end{aligned}$$

It follows that, for large S, optimal portfolio choice implies

$$\begin{aligned} E_t\left[\psi + E_{t+1}\left[u'(C_{t+1+S})\right]\right]R^f &= E_t\left[\left(\psi + E_{t+1}\left[u'(C_{t+1+S})\right]\right)R_{i,t+1}\right] \\ E\left[\psi + u'(C_{t+1+S})\right]R^f &= E\left[\left(\psi + u'(C_{t+1+S})\right)R_{i,t+1}\right] \\ &= Cov\left[u'(C_{t+1+S}), R_{i,t+1}\right] + E\left[\psi + u'(C_{t+1+S})\right]E\left[R_{i,t+1}\right] \end{aligned}$$

[6] The usual assumption is that the $\mu_t u'(C_t) - E_{t-1}[\mu_t u'(C_t)]$ is a covariance stationary process. Since we assume μ_t is stationary, joint stationarity is not particularly restrictive.

[7] See for example Aschauer (1985) Eichenbaum, Hansen, and Singleton (1988), Startz (1989), Flavin (2001), and Piazzesi, Schneider, and Tuzel (2003). It is uncertain but possible that ultimate consumption risk might also capture consumption risk in models with transitory movements in individual consumption risk or the cross-sectional distribution of marginal utilities such as the model of Lustig (2001).

so that we can apply a derivation parallel to that of the previous section to yield

$$\begin{aligned} E\left[R_{i,t+1}\right] &= R^f + \frac{Cov\left[\left(\ln\left(\frac{C_{t+1+S}}{C_t}\right), R_{i,t+1}\right]\right.}{Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]} \frac{\gamma Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}{E\left[1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]+\psi} \\ &= R^f + \beta_{i,S}\tilde{\lambda}_S \end{aligned}$$

where

$$\tilde{\lambda}_S = \frac{\gamma Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}{E\left[1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]+\psi} \tag{3.6}$$

Thus, for this class of models, we have the same structural beta representation, so that ultimate consumption risk should explain the cross-sectional pattern of returns. However, since $\tilde{\lambda}_S$ includes an additional term, ψ, in the denominator, our estimates of risk aversion are biased. A plausible assumption is that the stochastic discount factor based only on consumption understates the volatility of the true stochastic discount factor, and therefore that $Cov\left[\mu_t, u'\left(C_t\right)\right] = \psi > 0$. If this were true, our estimates of risk aversion that ignore the term ψ would be downward-biased. While there is no hard evidence on the magnitude of ψ, we suspect that, relative to the size of the denominator in equation (3.6), this bias is small.[8]

Related to this class of models, the appendix uses a different argument to show that there are a set of models with multiplicative external habits that do not have this bias in estimates of risk aversion. For this class of models, which includes that of Abel (1990), our measure of ultimate consumption risk prices average returns and our measure of risk aversion is unbiased, for S large enough.

Third, the ultimate risk also measures the risk of stocks in some models in which households face constraints on information, calculation, or adjustment of consumption so that consumption and marginal utility move only slowly to the new optimal level following a shock. As a concrete example, Gabaix and Laibson (2001) add to the canonical model of Merton (1969) the assumption that households face costs of monitoring their portfolio balances and so check and learn their wealth only infrequently, once every D periods.[9] When a household learns its wealth, it adjusts its consumption in response to all the market returns during the interval since it last learned its account balance. Assuming that a constant measure of households learn their balances and adjust their consumption at every instant, aggregate consumption adjusts smoothly and slowly over D periods to reflect a given return. In this case, the risk of a given portfolio under this model of economic behavior is measured by the ultimate risk to consumption.

[8] For example, Piazzesi, Schneider, and Tuzel (2003), consider μ_t to be the budget share of consumption spend on housing services (relative to its mean). Since both this and consumption fluctuate only by a few percent from their expected paths, $\psi << 1$.

[9] See also Caballero (1995), Lynch (1996), Marshall and Parekh (1999), Alvarez, Atkeson, and Kehoe (2000), and Sims (2001).

A final reason to evaluate the risk of equity with the ultimate risk is that aggregate consumption data may measure consumption responses with delay, even if the true consumption response were instantaneous. As demonstrated by Wilcox (1992), serially correlated measurement error is induced in aggregate consumption data by sampling error, imputation procedures, and definitional difficulties involved in constructing measures of aggregate consumption from monthly survey data on retail sales. For example, sales at a given retailer are allocated across types of consumption using shares that remain fixed for years. Given divergent price indexes, this induces serially correlated measurement error, only corrected once the correct shares are employed.

4. Data

For our portfolios and returns, we use quarterly returns on the Fama and French portfolios (25), which are the intersections of 5 portfolios formed on size (market equity, ME) and 5 portfolios formed on the ratio of book equity to market equity (B/M). B/M used during a fiscal year is based on the book equity for the previous fiscal year divided by ME for December of the previous year. The B/M breakpoints are the NYSE quintiles. The portfolios include all NYSE, AMEX, and NASDAQ stocks for which there is market equity data for December and June of the previous fiscal year, and (positive) book equity data for the previous fiscal year. We denote a portfolio by the rank of its ME and then the rank of its B/M so that the portfolio 15 is the smallest quintiles of stocks by ME and the largest quintile of stocks by B/M. The series are available monthly from July 1926 to December 2001.[10] To match consumption data, we use a quarterly frequency, and set our timing convention so that $R_{i,t+1}$ represents the return on portfolio i during the quarter $t+1$.

We measure consumption as personal consumption expenditures on nondurable goods from the National Income and Product Accounts, divided by the civilian non-institutional population of the United States. The data are made real using a chain-weighted price deflator, spliced across periods, produced by the Bureau of Economic Analysis. Except where noted, we make the "end of period" timing assumption that consumption during quarter t takes place at the end of the quarter. Our measure of the risk-free rate is the return on a three-month treasury bill. All returns are deflated by the same deflator as consumption.

These series determine our sample and frequency. We use a sample of returns for quarterly data from the second quarter of 1947 to the first quarter of 1998. We stop the sample of returns so that we can allow up to 4 years of consumption growth matched to a return ($S = 15$) without altering the sample of returns that we study as we vary S. That is, the sample for per capita consumption runs from the first quarter of 1947 to the last quarter of 1998 plus S quarters, which is the last quarter of 2001 when $S = 15$.

[10] We thank Kenneth French for making the Fama and French portfolio data available on his web page: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

5. Estimation and inference

Following the methodology of Fama and MacBeth (1973) and Fama and French (1992), we estimate equations (3.2) and (3.3) in two steps. First, each $\beta_{i,S}$ is estimated directly from the empirical counterparts to the theoretical moments in equation (3.3). This is simple to implement since $\hat{\beta}_{i,S}$ is the slope coefficient from a time series regression of return i onto a constant and $\ln \frac{C_{t+1+S}}{C_t}$.

Second, λ_S and its standard error are estimated from a series of cross-sectional regressions of returns on the $\hat{\beta}_{i,S}$, the slope coefficients estimated in the time series regressions. Let the parameters be denoted by $\theta_S = (\alpha_S, \lambda_S)'$, let the data be denoted by $R_{t+1} = (R_{1,t+1}, ...R_{25,t+1})'$ and $\hat{X}_S = \left(\iota, \hat{B}_S\right)$ where ι is a column vector of ones and $\hat{B}_S = \left(\hat{\beta}_{1,S}, ..., \hat{\beta}_{25,S}\right)'$. Let $\widehat{\theta}_{S,t}$ be the average of the estimates of θ_S from cross-sectional regression using only returns during period t

$$\widehat{\theta_S} = \frac{1}{T}\sum_{t=1}^{T}\widehat{\theta}_{S,t} \ , \ \ \widehat{\theta}_{S,t} = \left(\hat{X}_S'\hat{X}_S\right)^{-1}\hat{X}_S'R_t. \tag{5.1}$$

We calculate Fama and MacBeth standard errors from the observed variation in parameter estimates across subsamples,

$$\widehat{Var}\left(\sqrt{T}\left(\hat{\theta}_S - \theta_S\right)\right) = \sum_{t=1}^{T}\left(\widehat{\theta}_{S,t} - \tfrac{1}{T}\textstyle\sum_{t=1}^{T}\widehat{\theta}_{S,t}\right)\left(\widehat{\theta}_{S,t} - \tfrac{1}{T}\textstyle\sum_{t=1}^{T}\widehat{\theta}_{S,t}\right)'. \tag{5.2}$$

If the innovations to returns conditional on the factor are independently distributed over time then

$$\text{plim}\ \widehat{Var}\left(\sqrt{T}\left(\hat{\theta}_S - \theta_S\right)\right) = V = \left(X_S'X_S\right)^{-1}X_S'\Phi X_S\left(X_S'X_S\right)^{-1}$$

where $\Phi = E\left[\left(R_{t+1} - E\left[R_{t+1}\right]\right)\left(R_{t+1} - E\left[R_{t+1}\right]\right)'\right]$. However, these estimates of statistical uncertainty do not account for the fact that the betas are estimated regressors. Shanken (1992) shows that, if asset returns are homoskedastic and have a jointly normal distribution conditionally on the factor, then $\sqrt{T}\left(\widehat{\theta} - \theta\right)$ converges to a normal distribution with mean zero and variance $W = V + Q$ where Q is given by

$$Q = \lambda_S^2 Var\left[\ln\left(\tfrac{C_{t+1+S}}{C_t}\right)\right]^{-1}\left(X_S'X_S\right)^{-1}X_S'\Sigma X_S\left(X_S'X_S\right)^{-1} \tag{5.3}$$

where Σ is the constant covariance matrix of the residuals of the time series regressions. Since Q is positive definite, in this case, the Fama and MacBeth standard errors overstate the precision of the estimates. If the returns are indeed conditionally homoskedastic the Shanken standard errors are to be preferred. However, Jagannathan and Wang (1998) shows this ordering may not hold when the returns are conditionally heteroskedastic, leading us to report both standard errors. The Shanken standard error is estimated from the empirical counterparts to the elements of equation (5.3).

Turning from the statistical importance of consumption risk to the economic importance of consumption risk, we report a cross-sectional R^2 statistic from our second-stage regressions. Following Jagannathan and Wang (1998), the cross-sectional R^2 measures the fraction of cross-sectional variation explained by the data given the average θ from the cross-sectional regressions,

$$1 - Var_c\left(\frac{1}{T}\sum_{t=1}^{T} R_{i,t+1} - \hat{\alpha}_S - \hat{\lambda}_S \hat{\beta}_S^i\right) / Var_c\left(\frac{1}{T}\sum_{t=1}^{T} R_{i,t+1}\right),$$

where $Var_c(.)$ denotes a cross-sectional variance. Since the $\hat{\beta}_{i,S}$ are estimated, it is also the case that the fit of the second stage regression is biased downwards in finite samples. To gauge the importance of this bias, we estimate the noise and signal in the $\hat{\beta}_{i,S}$ for different S. This also allows us to evaluate the relative importance of the added signal (additional movement in expected consumption growth) and added noise (innovations to consumption after time $t+1$) in the $\hat{\beta}_{i,S}$ as we increase S.

We report inference concerning the implied risk-free rate and risk aversion based only on the statistical uncertainty of the coefficients in the cross-sectional regression. That is, the standard error of the risk-free rate is the standard error of α, and the test of its equality to the observed average risk-free rate takes the empirical moment, $\hat{E}\left[R^f\right]$, as given. Similarly, we make inference regarding the risk aversion of the representative investor taking as given the observed values of $Var\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]$ and $E\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]$. This approach isolates the uncertainty about risk aversion due to the relationship between consumption risk and average returns, but understates our true degree of ignorance about risk aversion and the risk-free rate.

6. Main findings

This section asks whether consumption risk explains the cross-sectional variation in expected returns on different portfolios of stocks. First, do the $\hat{\beta}_{i,S}$ explain a large share of the variance of average returns – is consumption risk economically significant? Second, is $\hat{\lambda}_S$ significant – is consumption risk statistically significant? These two questions address the major issue of whether consumption risk is or is not a major determinant of the cross-sectional pattern of returns.

We pursue two additional checks of the model and, implicitly, of the importance of consumption risk. First, $\hat{\alpha}_S$ is an estimate of the implied risk-free rate. Thus, an additional test of the model is whether this estimate is similar to the sample average of the risk-free rate. Second, we recover an estimate of the risk aversion of the representative investor. Since this is a structural parameter that governs other household behaviors, a test of the model is whether the estimated risk aversion is plausible given what it would imply for economic behavior and risk taking in general.

6.1. Baseline results

We begin by reporting the economic and statistical significance of contemporaneous consumption risk for the returns of these portfolios. The first two rows of Table 1 show that for $S = 0$, consumption betas explain only 12 or 19 percent of the cross-sectional variance in average returns. The first row reports results using the traditional "start of period" timing convention for consumption growth that aligns $\Delta \ln C_{t+2}$ with $R_{i,t+1}$, while the second row reports results using the end of period timing convention that aligns aligns $\Delta \ln C_{t+1}$ with $R_{i,t+1}$.[11] For $S > 0$, we report results using the end of period timing convention since only for this convention is the initial period entirely prior to the period covered by the return.

In addition to having a low fit, contemporaneous risk is not statistically significant in explaining average returns. Both the Fama-MacBeth and Shanken standard errors are of similar magnitude to the estimated slope coefficient $\hat{\lambda}$. The insignificant slope coefficient is accompanied by a large estimate of the intercept, which implies a large estimate of the risk-free rate. The fifth and sixth entries in the first two rows of of Table 1 show that the annual real interest rate is estimated to be 15.0 or 8.6 percent, far in excess of and statistically significantly different from the actual average risk-free rate over the sample of 1.7 percent. Finally, the low slope coefficient and high intercept imply implausibly large levels of household risk aversion. The point estimate of the risk aversion required of the representative agent to rationalize the spread in average returns given the differences in the contemporaneous consumption risk of the returns is 41 or 55. A 95 percent confidence interval contains reasonable levels of risk aversion only because statistical uncertainty is large.

The remainder of Table 1 shows that, contrary to the result for contemporaneous risk, consumption risk measured after consumption has had time to adjusts to returns explains a significant share of the variance in average returns. As we increase S, the inferred economic importance of consumption risk rises: consumption betas measured using the ultimate risk explain as much as half the variance in average returns across portfolios. This rise is not monotonic, and the decline for $S = 12$ and higher is discussed subsequently.

The magnitude of the slope coefficient, $\hat{\lambda}$, rises monotonically with S, and is statistically significant for horizons 9 and higher. The slope coefficient measures the extent to which one portfolio is expected to have a higher return than another, based on the difference in their risk, as measured by the difference between the covariance of the returns on each portfolio and the following change in consumption over $1 + S$ periods. While this evidence seems strong, if the model is misspecified, it is possible for the t-statistic to not converge to zero (Kan and Zhang (1999)). We provide a test of misspecification later in this section of the paper. For now, for large S, we observe that

[11] Campbell (1999) uses the start of period convention because, consistent with the finding here with respect to the cross-section of returns, the start of period convention produces a larger contemporaneous covariance between consumption and aggregate market returns than the end of period timing convention.

consumption risk explains average returns.

The greater economic and statistical importance of consumption risk for large S is accompanied by an increased plausibility of the implied interest rate on the risk-free asset and the implied risk aversion of the representative agent. For $S > 6$ the average risk-free real interest rate and the intercept of the cross-sectional regression are statistically indistinguishable, although the gap between the estimate and the actual average remains $2 - 5$ percent. The implied risk aversion of the representative agent declines nearly monotonically in S to level off at a point estimate of around 12. The statistical uncertainty declines substantially with S, so that for large S, a 95 percent confidence interval contains roughly any positive number below 15. The 95 percent confidence interval includes reasonable estimates of the representative agent's risk aversion not simply because of large statistical uncertainty.

But how should we choose the correct horizon of consumption adjustment, S? What are our best estimates for fit and parameter estimates given that they vary by horizon? On the one hand, if consumption is slow to adjust for the reasons outlined above, then larger S are preferred since they allow a longer horizon for consumption adjustment. The fact that the estimates of risk aversion and the risk-free rate change with S are evidence against the textbook C-CAPM and in favor of the C-CAPM with slow consumption adjustment. This suggests that we would like to choose S as large as possible. On the other hand, as one increases S, one also increases the noise in the estimates of β^i_S from the time series regressions. Measurement error in a factor, here $\ln\left(\frac{C_{t+1+S}}{C_t}\right)$, is irrelevant asymptotically: the effect on the denominator of each β^i_S is offset by a corresponding change in the numerator of λ. But measurement error decreases the accuracy of estimation of β^i_S in any finite sample. For larger S, measurement error is greater because we do not observe the expectation at time $t+1$ of consumption growth over the subsequent S periods, but instead must construct or proxy it.[12] Closely related, for larger S, there is more correlation over time in residuals in the time series regressions that estimate each β^i_S and thus less precisely estimated β^i_S.[13] These arguments imply that we do not want to choose S too large.

One solution would be to use the value of S which gives the best estimates of the structural parameters. But such an approach would undermine inference. Instead we consider two different approaches to choosing a "best" value for S.

First, we choose S based on the statistical uncertainty in the β^i_S's from the time series regressions relative to the total cross-sectional variance of the $\hat{\beta}^i_S$ used in the

[12] We use actual consumption growth in place of consumption growth expected given information at $t+1$. An alternative would be to use a constructed measure of $E_{t+1}\left[\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]$. But this measure would still contain noise, and potentially more importantly would omit signal – true slow adjustment of consumption – from our factor.

[13] There is also a reason to keep the maximum S small and not include larger values of S in our reported results. A larger maximum S studied implies a shorter time series of available returns on which to estimate all β_S, so that imprecision in all $\hat{\beta}_S$ increases.

second stage. We calculate for different values of S the ratio

$$\frac{\frac{1}{25}\sum_{i=1}^{25}\widehat{Var}\left(\hat{\beta}_S^i\right)}{\frac{1}{25}\sum_{i=1}^{25}\left(\hat{\beta}_S^i-\frac{1}{25}\sum_{i=1}^{25}\hat{\beta}_S^i\right)^2}$$

where the numerator is the variance of $\hat{\beta}_S^i$ in each time-series regression averaged across portfolios and the denominator is the cross-sectional variance in $\hat{\beta}_S^i$. The numerator is a measure of the average noise in $\hat{\beta}_S^i$ and the denominator is a measures of the total observed signal plus noise across β_S^i. The calculation of the numerator is based upon standard errors calculated accounting for time-series correlations using Newey-West standard errors in each time-series regression. Because any correlation among $\hat{\beta}_S^i$ affects both numerator and denominator, the measure is not perfect and the numerator can be greater than the denominator, and indeed is. Figure 1 shows that statistical uncertainty in the estimated β_S^i remains a roughly constant share of the total variance from $S = 0$ to 11 and rises significantly after this. On these grounds, we choose $S = 11$, since greater values of S lead to much larger ratios of statistical uncertainty to observed variation in consumption beta's. Consistent both with this evidence and the hypotheses that consumption risk is important and that consumption adjusts slowly, Table 1 shows that $S = 11$ is also the horizon for which the consumption betas explain the largest fraction of average returns and our estimates of risk aversion are most precise. On the other hand, still larger values of S lead to greater precision in estimates of λ and more realistic estimates of the risk-free interest rate and the coefficient of relative risk aversion.

Second, consider the choice of S as a model selection problem in which S indexes a set of models none of which are true and in which the sample of data is given.[14] As discussed by Leamer (1983), the best model minimizes the difference between the true and approximating density for the data in the information sense. Thus, choosing among the models amounts to choosing the model which produces the best fit, since there is no difference in the number of parameters across models. Thus, again we are lead to choose $S = 11$.

In our tables, we continue to report all horizons, and to report in bold the model selected by our first criterion, the variance ratio. We focus on the first criterion since it does not necessarily select the model with the best fit. In all specifications and samples that we have analyzed, the exact choice of S does not drive our inference on structural parameters, in that estimates are similar for models (S) near the selected model.

Returning to our main results, Figure 2 plots the consumption betas, average realized returns, and the second-stage regression lines for $S = 0$ and for $S = 11$. One

[14] If the sample were to grow, we could reasonably think we were estimating the true model by choosing S large enough (or letting S grow with the sample) and our estimates would be consistent, although potentially uninformative in our sample. Since instead we are fixing the sample and treating mismeasurement of β as a fixed problem for inference, we are naturally lead to Bayesin methods.

can see that for $S = 11$, the $\hat{\beta}^i$ explain more of the variance in returns. If the model is correct, as the estimated expected returns and $\hat{\beta}^i$ converge to their true values, the fit of the second stage regression should become perfect. One can also observe that the lower estimate of the risk-free rate is due to the increased slope of the regression line, although this is not a precise statement since the scale of the β's is different due to the different variance of the stochastic discount factor.

Figure 3 recasts Figure 2 in terms of the more economically interesting concept of pricing errors. The horizontal axis displays average returns and the vertical axis predicted returns based on consumption betas measured for $S = 0$ and for $S = 11$. Each portfolio is denoted by the relative size (1 to 5) and relative book to market ratio (1 to 5). The error in pricing a return using consumption risk is the vertical difference between the fitted return and average return and given by the distance of a point from the 45-degree line. For all but five of the twenty five portfolios (ME 1, B/M 3; ME 3, B/M 2; ME 4, B/M 1; ME 5, B/M 3; and ME 5, B/M 5), the pricing error is smaller for $S = 11$ than for $S = 0$. Table 2 presents the information slightly differently, showing the square root of the average squared pricing error for each way of splitting into five quintiles. The pricing errors are aggregated to five size and then five book to market quintiles. Comparing horizon $S = 11$ to horizon $S = 0$, the pricing errors are lower in four out of five of size quintiles and four out of five book equity to market equity quintiles. In each of the two exceptions (the middle book to market portfolio and the largest size portfolio), the pricing error associated with contemporaneous consumption risk, $S = 0$, is the smallest of the portfolios. Further, the increases in the pricing error across horizons for these two portfolios are smaller than the decreases in pricing error for every one of the remaining portfolios.

Before turning to comparisons with factor models of returns, we conclude this section of the paper by reporting the results of estimating our model in several different ways and the results of testing for misspecification using portfolio characteristics.

6.2. Alternative specifications

We consider three alternative ways of estimating of our model. First, previous work has focussed on a shorter time period than we analyze in our baseline results (e.g. Fama and French (1992 and 1993)). Table 3 presents results for a sample set to match that of Lettau and Ludvigson (2001b). This sample uses returns from the third quarter of 1963 to the third quarter of 1998, which implies that we can only study a maximum horizon of three and a half years (S= 13). In this sub-period, our slow consumption adjustment model does even better at explaining average returns. The data still suggest that $S = 11$ is the best horizon to study, and at this horizon consumption risk explains 73 percent of the variation in average returns. The implied risk-free rate is below the average return on a three-month treasury bill. On the other hand, estimated risk aversion rises slightly to 14.

Second, Ait-Sahalia, Parker, and Yogo (2001) argue that the consumption risk of equity is understated by NIPA nondurable goods because it contains many necessities and few luxury goods. Parker (2001) shows that this concern can be partly addressed when using long-term measures of consumption risk by conducting the analysis with total consumption instead of nondurable consumption. The usual concern with using total consumption is that it contains expenditures on durable goods instead of the theoretically desired stock of durable goods. But expenditures and stocks are cointegrated. The long-term movement in expenditures following an innovation to equity returns also measures the long-term movement in consumption flows.

Table 4 presents our main results for total consumption per capita. Compared to the results for nondurable goods, the initial fit for $S = 1$ is better and the increase in fit for larger S is not monotone, but declines over the first year. As with the main results, the noise ratio suggests that $S = 11$ provides the best estimates as the noise to cross-sectional variance ratio rises significantly at $S = 12$ and each larger S. At $S = 11$, the fit is higher than for nondurable consumption, 56 percent, and the estimated coefficient of relative risk aversion is lower, below 10. On the other hand, the implied risk-free rate, while still statistically indistinguishable from the average treasury bill rate, is economically farther from it, only slightly below 6 percent.

Third, we consider a slightly different set of returns, the equal-weighted Fama and French (25) portfolios. Table 5 shows that ultimate consumption risk does an even better job of explaining the cross-sectional pattern of returns of these (albeit similar) portfolios. The estimated risk-free return is 3.0 percent as opposed to 5.3 percent for value-weighted portfolios. On the other hand, the estimated level of risk aversion is slightly higher using this set of portfolios.

6.3. Testing the model for misspecification

While the evidence presented so far is consistent with an important role for consumption risk, if our model were misspecified, the cross-sectional estimation could be asymptotically biased. Kan and Zhang (1999), with a simulation exercise, show that "useless" factors can appear statistically significant when the Fama and MacBeth methodology is applied to a misspecified model. When the factors are misspecified, the t statistic for a beta may converge in probability to infinity even when the true coefficient is zero. Therefore, it is possible that the statistical significance of the ultimate risk to consumption might be spurious.

To test the hypothesis of misspecification, we follow Jagannathan and Wang (1996) and expand our model to include other asset characteristics such as book-to-market ratio and relative market size. Adding these to equation (3.2), we obtain

$$E\left[R_{i,t+1}\right] = a_S + a_z' z_i + \beta_{i,S}\lambda_S \tag{6.1}$$

where z_i is the time average of the vector of observable asset characteristics. Using formulae analogous to those in Section 5, we construct t-statistics for the coefficient on the asset characteristics using both the Fama and MacBeth standard errors and the standard errors with the Shanken correction. If the model in equation (3.2) is correctly specified, then a_z should be a vector of zeros. If instead the model is misspecified, the

t-statistics for firm characteristics converge to infinity in probability (Jagannathan and Wang (1998)). Thus we include asset characteristic in our cross-sectional regressions to detect model misspecification, since large t-statistics for firm characteristics would reject the beta pricing model under analysis.

Table 6 shows the fit of the regressions of returns onto the ultimate consumption risk betas and asset characteristics for two different asset characteristics – size and book to market value – and two different measures of each of these characteristics – the mean of the log of the characteristic over time for each portfolio or the log of the mean of the characteristics over time for each portfolio. First note that in all four cases, although only trivially for the log of the mean book to market value, the fit of the model still is larger for the horizon $S = 11$ than for the horizon $S = 0$. Second, in all four cases the statistical significance of the coefficient on consumption beta is larger for $S = 11$ than for $S = 0$. Thus, slow adjustment of consumption adds explanatory power to consumption risk even in the presence of firm characteristics. For all panels except the upper left panel, the statistical significance of firm characteristics is lower for $S = 11$ than for $S = 0$. Finally, and most relevant for the test at hand, only for one panel is the firm characteristic significant at a horizon of 11, while only for one panel is consumption risk not significant at a horizon of 11.

Having shown that the ultimate risk to consumption is a significant explanator of average returns across portfolios of stocks, we now compare the performance of consumption risk as a "factor" to two important extant factor models that have been used to study the average returns in the Fama and French portfolios. A final section of results follows in which we derive and analyze the ultimate risk to consumption from the textbook C-CAPM in an economy with a time-varying risk-free rate of interest.

7. Relation to Fama-French and Lettau-Ludvigson factor models

Fama and French (1992) and Fama and French (1993) show that three factors explain a large fraction of the cross-sectional variation in returns in the Fama and French portfolio: the overall market return (denoted R^m), the difference between the size of firms in the smallest and largest size quintiles ("small minus big" denoted SMB), and the difference between the book to market equity ratios in the largest and smallest book to market equity quintiles ("high minus low" denoted HML). Lettau and Ludvigson (2001a) argues that the budget constraint of the representative household implies that consumption, income and wealth should be cointegrated and then shows that the deviation of these variables from their long-run relationship (the error-correction term in the three variable vector autoregression) is a good predictor of market returns. Lettau and Ludvigson (2001b) shows that this variable, denoted by cay_t, consumption growth ($\Delta \ln C_{t+1}$), and their interaction provide a three factor model that does as well in explaining the average cross-section of returns as the Fama and French three factor model. In this section, we compare the fit of our one factor model to these alternative factor models.

Both three factor models assume that there exists a stochastic discount factor that prices returns, M_{t+1}, which is a linear function of the factors, which we denote f_{t+1}. In the Lettau and Ludvigson model for example, we have

$$M_{t+1} = b_0 + b_1 cay_t + b_2 \Delta \ln C_{t+1} + b_3 cay_t \Delta \ln C_{t+1}.$$

This model has the beta representation

$$E\left[R^i_{t+1}\right] = \alpha + \beta_i' \lambda \tag{7.1}$$

where

$$\alpha = R^f,\ \beta_i = Var\left[\mathrm{f}_{t+1}\right]^{-1} Cov\left[\mathrm{f}_{t+1}, R_{i,t+1}\right],\ \lambda = -R^f Var\left[\mathrm{f}_{t+1}\right] \mathrm{b}$$

and $\mathrm{f}_{t+1} = \left(cay_t, \Delta \ln C_{t+1}, cay_t \Delta \ln C_{t+1}\right)'$, the vector of factors, and b$= (b_1, b_2, b_3)'$ the vector of coefficients on the factors in the stochastic discount factor. In the Fama and French three factor model, $\mathrm{f}_{t+1} = \left(R^m_{t+1}, SMB_{t+1}, HML_{t+1}\right)'$. We estimate these models using the Fama-MacBeth methodology, first by estimating the vector β_i from a series of multivariate time-series regressions of the factors on the returns and then by estimating equation (7.1) on the cross-section of $\hat{\beta}_i$'s. We report both standard errors from the Fama-MacBeth procedure and those with the Shanken adjustment.

We initially focus on the subsample analyzed by Lettau and Ludvigson (2001b) and the focus of Fama and French (1993). The first row of results in the first panel of Table 7 reports the fit, implied risk-free rate, and coefficients on the Fama and French factors and their standard errors, the Fama-MacBeth standard error above the Shanken-adjusted standard error. The penultimate column reports the t-tests of the hypothesis that the risk-free rate implied by the model equals the average Treasury bill return in the sample. The second row of results reports the same set of statistics for the Lettau and Ludvigson three-factor model (using our timing convention for consumption). One interpretation of this model is that the coefficient of relative risk aversion is time-varying, so that the stochastic discount factor is $1 - \gamma_t \Delta \ln C_{t+1}$.[15] Since risk aversion appears only multiplied by consumption growth, we also report in the third row results which omit cay_t on its own. We view this as consistent with this more structural view of the role of cay_t, that it captures time variation in risk aversion.[16] The final column of results shows the point estimate and a 95-percent confidence interval for the average level of risk aversion implied by the estimated

[15] We assume that marginal utility at $t + 1$ is given by $C_{t+1}^{-\gamma_t}$ and make the arbitrage relation stationary by dividing by $C_t^{-\gamma_t}$.

[16] There are two possible interpretations of the role of *cay*. Time-varying risk aversion is one; the other, less structural, is that *cay* is a better version of a dividend-price ratio, where dividend is measured by c and y and price by a. Lettau and Ludvigson (2001) state "some stocks tend to be more highly correlated with consumption growth in bad times, when **risk or risk aversion** is high, than they are in good times, when **risk or risk aversion** is low [emphasis added]."

coefficients in this model.

$$\gamma = -\left(b_2 + b_3 \frac{1}{T}\sum_{t=1}^{T} cay_t\right)$$

The fourth row presents the results for our single-factor model, the ultimate risk to consumption, and the results correspond to those in Table 3 with a horizon of 11 quarters. The final row reports the results for the textbook C-CAPM and corresponds to Table 3 row 2.[17]

The first main point of Table 7 is that the fit of the Fama and French, Lettau and Ludvigson, and ultimate risk to consumption models are all good: 78 percent, 66 percent and 73 percent of the variation in returns are explained by each model's betas respectively. The horizon $S = 11$ consumption beta is a single factor model, and yet it explains returns as well as these three-factor models. Second, while the factors are jointly significant, the individual factors in both three-factor models are all statistically insignificant except HML. The ultimate risk to consumption is highly statistically significant. Third, the betas on the Fama and French factors and the beta on the ultimate risk to consumption both imply reasonable levels of the risk-free rate of interest, while the Lettau and Ludvigson model performs less well on this dimension. Fourth, a fair amount of explanatory power in the Lettau and Ludvigson model is lost when the factor *cay* alone is removed. Consumption growth interacted with *cay* to capture time-variation in risk aversion still explains more of the variation in returns than the C-CAPM, but also now noticeably less than the Fama and French model and ultimate consumption risk. Related to this point, the magnitude of the estimates of risk aversion from the Lettau and Ludvigson model are much higher than for the ultimate risk to consumption and suggest that *cay* does not measure time variation in risk aversion. Finally, and most well-known, the textbook C-CAPM performs poorly.

Figure 4 graphs the pricing errors for each portfolio as the difference on the vertical axis between a portfolio and the 45 degree line, for the four main models. All models besides the textbook C-CAPM do quite well at fitting returns, although in this sample and with this timing assumption, the performance of the textbook C-CAPM is not terrible.

The second panel of Table 7 reports the same set of information as the first panel but expands the sample to the longest possible sample and thus the closest to the baseline sample used initially in this paper.[18] The same conclusions stand as for the shorter sample. Figure 5 plots fitted and average returns in this sample, but plots the Lettau and Ludvigson two factor model that omits the variable *cay* not interacted with consumption growth.

[17] The textbook C-CAPM performs significantly worse than reported here under the start-of-period timing convention.

[18] The starting date is determined by the first available observation on cay_t.

8. Consumption risk when the risk-free rate varies over time

So far we have assumed that the risk-free rate in the economy is constant. In this section we demonstrate that our main results do not depend on this approximation. We consider the textbook model only. When the risk-free rate varies, the extent to which the model based only on the ultimate risk as in Section 3 differs from the truth is related to the extent to which an innovation to returns leads to a change in future risk-free rates. Intuitively, according to the textbook model, if an innovation to returns were to lead to a significant revision in planned intertemporal substitution in consumption over the next S periods, then looking S periods out could be quite misleading. If the only reason for consumption to move from planned consumption between periods $t+1$ and $t+1+S$ were future innovations, then this measure would exactly measure consumption risk. Thus, we demonstrate that our findings remain when we include as an additional factor consumption growth interacted with changes in the risk-free rate.

8.1. The ultimate risk to consumption

In an economy with a time-varying real interest rate, the textbook household optimization problem implies that consumption obeys an intertemporal Euler equation between $t+1$ and $t+1+S$ for the risk-free rate

$$E_{t+1}\left[\frac{\delta^S R^f_{t+1,t+1+S} u'\left(C_{t+1+S}\right)}{u'\left(C_{t+1}\right)}\right] = 1.$$

This Euler equation implies that marginal utility in $t+1+S$ is equal to (discounted) marginal utility in $t+1$ plus an innovation to marginal utility due to information that arises between $t+1$ and $t+1+S$, denoted $\varepsilon_{t+1,t+1+S}$,

$$u'\left(C_{t+1}\right) = \delta^S R^f_{t+1,t+1+S} u'\left(C_{t+1+S}\right) - \varepsilon_{t+1,t+1+S}.$$

Substituting this relationship into the arbitrage equation (2.1), gives

$$E_t\left[R^f_{t+1,t+1+S} u'\left(C_{t+1+S}\right) R_{i,t+1}\right] - E_t\left[R^f_{t+1,t+1+S} u'\left(C_{t+1+S}\right)\right] R^f_{t,t+1} = 0.$$

Following the same derivation as before leads to a beta representation with

$$\begin{aligned}
\alpha_S &= E\left[R^f_{t,t+1}\right] + \frac{Cov\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}, R^f_{t,t+1}\right]}{E\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}\right]}, \\
\beta_{i,S} &= \frac{Cov\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}, R_{i,t+1}\right]}{Var\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}\right]}, \\
\lambda_S &= -\frac{Var\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}\right]}{E\left[R^f_{t+1,t+1+S}\frac{u'(C_{t+1+S})}{u'(C_t)}\right]}.
\end{aligned}$$

Log-linearizing consumption growth, $\frac{u'(C_{t+1+S})}{u'(C_t)} \approxeq 1 - \gamma \ln\left(\frac{C_{t+1+S}}{C_t}\right)$, we obtain a two factor beta representation

$$E\left[R^i_{t,t+1}\right] = \alpha_S + \beta^C_{i,S}\lambda^C_S + \beta^f_{i,S}\lambda^f_S \tag{8.1}$$

where:

$$\alpha_S = E\left[R^f_{t,t+1}\right] + \frac{Cov\left[R^f_{t+1,t+1+S}\left(1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right), R^f_{t,t+1}\right]}{E\left[R^f_{t+1,t+1+S}\left(1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right)\right]};$$

$$\beta^f_{i,S} = \frac{Cov\left[R^f_{t+1,t+1+S}, R^i_{t,t+1}\right]}{Var\left[R^f_{t+1,t+1+S}\right]};\ \lambda^f_S = -\frac{Var\left[R^f_{t+1,t+1+S}\right]}{E\left[R^f_{t+1,t+1+S}\left(1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right)\right]};$$

$$\beta^C_{i,S} = \frac{Cov\left[R^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right), R_{i,t+1}\right]}{Var\left[R^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]};\ \lambda^C_S = \frac{\gamma Var\left[R^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}{E\left[R^f_{t+1,t+1+S}\left(1-\gamma\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right)\right]}.$$

The two factors are consumption growth interacted with the future risk-free rate ($R^f_{t+1,t+1+S}$) and the future risk free rate alone. Thus we have a "scaled" linear factor model in which $R^f_{t+1,t+1+S}$ is the scaling factor. Note that if we assume that $R^f_{t+1,t+1+S}$ and $R^i_{t,t+1}$ are orthogonal, then we have that $\beta^f_{i,S}\lambda^f_S = 0$ and our model collapses to a one-factor model with $\beta^C_{i,S}$ only.

We also recover a measure of the implied risk aversion from both $\widehat{\lambda}^C_S$ and $\widehat{\lambda}^f_S$. These measures are, respectively,

$$\hat{\gamma}^C_S = \frac{\hat{\lambda}^C_S E\left[\widehat{R}^f_{t+1,t+1+S}\right]}{Var\left[\widehat{R}^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right] + \hat{\lambda}^C_S E\left[\widehat{R}^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}, \tag{8.2}$$

$$\hat{\gamma}^f_S = \frac{\widehat{\lambda}^f_S E\left[\widehat{R}^f_{t+1,t+1+S}\right] + Var\left[\widehat{R}^f_{t+1,t+1+S}\right]}{E\left[\widehat{R}^f_{t+1,t+1+S}\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right]}. \tag{8.3}$$

Using the estimated relative risk aversion we can recover a measure of the implied mean risk free rate

$$\widehat{E}\left[R^f_{t,t+1}\right] = \widehat{\alpha}_S - \frac{Cov\left[\widehat{R}^f_{t+1,t+1+S}\left(1-\widehat{\gamma}_S\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right), R^f_{t,t+1}\right]}{E\left[\widehat{R}^f_{t+1,t+1+S}\left(1-\widehat{\gamma}_S\ln\left(\frac{C_{t+1+S}}{C_t}\right)\right)\right]}. \tag{8.4}$$

8.2. Estimation

Since $R^f_{t+1,t+1+S}$ is not observable for $S > 1$, the appropriate risk free rate need to be constructed. We approximate the risk-free rate between $t+1$and $t+1+S$ by the expected value of the product of the one-period risk-free rates

$$R^f_{t+1,t+1+S} \approx E_{t+1}\left[\prod_{j=0}^{S-1} R^f_{t+1+j,t+2+j}\right].$$

To model conditional expectations, we estimate the forecasting model

$$\prod_{j=0}^{S-1} R^f_{t+1+j,t+2+j} = \Phi_0 + \Phi_1 Z_t + \Phi_2 Z_{t-1} + ... + \Phi_{p+1} Z_{t-p} + \varepsilon_t \tag{8.5}$$

where Φ are parameter vectors and $Z_t = \left[R^f_{t,t+1},\, \pi_t,\, \Delta \ln C_t\right]'$. Our series for long-horizon risk-free rates is

$$\begin{aligned} R^f_{t+1,t+1+S} &= \hat{E}_{t+1}\left[\prod_{j=0}^{S-1} R^f_{t+1+j,t+2+j}\right] \\ &= \hat{\Phi}_0 + \hat{\Phi}_1 Z_t + \hat{\Phi}_2 Z_{t-1} + ... + \hat{\Phi}_{p+1} Z_{t-p} \end{aligned}$$

The fit in the forecasting regression exceeds 90 percent, suggesting that the assumed approximation is reasonable.

Estimation of the two factor model uses the Fama-MacBeth procedure. $\beta^C_{i,S}$ and $\beta^f_{i,S}$ are estimated by the time-series regressions

$$R_{i,t+1} = \alpha^f_S + \beta^f_{i,S} \widehat{R}^f_{t+1,t+1+S} + \epsilon^f_{S,t,t+1}, \tag{8.6}$$

$$R_{i,t+1} = \alpha^C_S + \beta^C_{i,S}\left[\widehat{R}^f_{t+1,t+1+S} \ln\left(\frac{C_{t+1+S}}{C_t}\right)\right] + \epsilon^C_{S,t,t+1}, \tag{8.7}$$

where we introduced the convention $\widehat{R}^f_{t+1,t+2} = R^f_{t+1,t+2}$ and $\widehat{R}^f_{t+1,t+1} = 1$. In the second stage, a series of cross-sectional regressions of equation (8.1) estimate λ^C_S, λ^f_S, and their statistical uncertainty.

8.3. Results

We estimate the model in equation (8.1) using the Fama-MacBeth procedure. The second column of Table 8 reports the cross-sectional R^2 of the two factor model given by equation (8.1) while the remaining columns report statistics and estimated parameters for the one factor model in which we include only the factor $R^f_{t+1,t+1+S} \ln\left(\frac{C_{t+1+S}}{C_t}\right)$

and not $R^f_{t+1,t+1+S}$ (the one factor model is the same as in equation (8.1) with the restriction $\beta^f_{i,S}\lambda^f_S = 0$). A comparison between the second and the third columns shows that adding $\widehat{R}^f_{t+1,t+1+S}$ as additional factor increases the average fit of the model by only 3 percent. Therefore, it is reasonable to concentrate on $\widehat{R}^f_{t+1,t+1+S} \ln\left(\frac{C_{t+1+S}}{C_t}\right)$ as the only factor. The first two rows of Table 8 reports the same results as the first two rows of Table 1 since, for $S = 0$, we have that $\widehat{R}^f_{t+1,t+1} = 1$ and $\beta^f_{i,S}\lambda^f_S = 0$. As S increases, the economic importance of the factor increases and for $S = 11$ it explains as much as 54 percent of the cross-sectional variance in average returns. As in Table 1, this rise is not monotone and it declines for $S \geq 12$. This corresponds to the starting point of a sharp rise in the ratio of statistical uncertainty to observed variation in $\hat{\beta}^C_{i,S}$. The cross-sectional R^2 of the one factor model is on average 1% higher than the one reported in Table 1, which shows that adding $\widehat{R}^f_{t+1,t+1+S}$ as a scaling variable does not decrease the high explanatory power of the ultimate risk to consumption, but actually increases it slightly.

The magnitude of the slope coefficient, $\widehat{\lambda}^C$, rises monotonically with S and is statistically significant for horizons 8 and higher. Thus, for large enough S, consumption risk measured by $\widehat{R}^f_{t+1,t+1+S} \ln\left(\frac{C_{t+1+S}}{C_t}\right)$ explains the cross-section of returns. Moreover, $\widehat{\lambda}^C$ is never statistically different from the point estimates of $\widehat{\lambda}$ reported in Table 1 and exhibits an almost identical pattern. We find similar conclusions for the implied risk-free rate and risk aversion. $\hat{\alpha}$ decreases almost monotonically in S and becomes statistically indistinguishable from the time series average of the three-month Treasury bill for $S > 6$. The implied risk aversion declines in S and levels off at a point estimate around 11. Furthermore, for larger S, the statistical uncertainty is much lower and a 95 percent confidence interval contains roughly all positive numbers below 14.

The analysis presented in this section is coherent, both qualitatively and quantitatively, with the results obtained under constant risk free rate assumption, and shows that the assumption of a constant risk-free rate or return is not important for our previous conclusions.

9. Conclusion

This paper analyzes the Fama and French portfolios and measures their riskiness by their ultimate risk to aggregate consumption. When investors are allocating their portfolios efficiently, differences in expected returns on assets should be explained by differences in the risk of each marginal investment to the utility of investors. We show that while the covariance of each portfolio and contemporaneous consumption growth has almost no predictive power for explaining the pattern of average returns across portfolios, the ultimate risk to consumption is highly statistically significant in explaining average returns and explains a large fraction of the variation in average returns. The average risk-free real rate of interest and the coefficient of relative risk aversion of the representative household calculated from the fitted model are more reasonable and more precise than those obtained in the previous literature; confidence intervals include completely reasonable estimates of these structural parameters. The fit of our one factor model rivals that of the three factor models of Fama and French and Lettau and Ludvigson. These conclusions are robust to several variations in assumptions.

In sum, the insights of the C-CAPM are alive and well: consumption risk is an important determinant of relative returns.

These results raise several questions. First, what is the true stochastic discount factor that prices returns, over time and across assets? Second, and more realistically, while the ultimate risk to aggregate consumption helps to explain the premium of equity over and above bonds, estimated risk aversion remains implausibly large. However, Parker (2001) narrows the focus to households that actually hold stock, and shows that the ultimate consumption risk of stockholders comes close to rationalizing the observed premium on equity. Is the consumption of stockholders alone an even better explanator of the cross-section of returns and does it imply more reasonable risk aversion in this cross-sectional context? A final open question is whether the ultimate risk to consumption explains differences in bond returns or differences in aggregate returns over time.

References

ABEL, A. B. (1990): "Asset Prices Under Habit Formation and Catching Up with the Joneses," American Economic Review, 80, 38–42.

AIT-SAHALIA, Y., J. A. PARKER, AND M. YOGO (2001): "Luxury Goods and the Equity Premium," NBER Working Paper 8417.

ALVAREZ, F., A. ATKESON, AND P. J. KEHOE (2000): "Money, Interest Rates, and Exchange Rates with Endogenously Segmented Asset Markets," NBER Working Paper 7871.

ASCHAUER, D. (1985): "Fiscal Policy and Aggregate Demand," American Economic Review, 75, 117–27.

ATTANASIO, O., AND G. WEBER (1993): "Consumption Growth, the Interest Rate and Aggregation," Review of Economic Studies, 60, 631–49.

——— (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," Journal of Political Economy, 103(6), 1121–57.

BANSAL, R., R. F. DITTMAR, AND C. T. LUNDBLAD (2001): "Consumption, Dividends, and the Cross-Section of Equity Returns," Manuscript, Fuqua School of Business, Duke University.

BASU, S., AND M. S. KIMBALL (2000): "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption," University of Michigan Manuscript.

BAXTER, M., AND U. J. JERMANN (1999): "Household Production and the Excess Sensitivity of Consumption to Current Income," American Economic Review, 89, 902–20.

BRAINARD, W. C., W. R. NELSON, AND M. D. SHAPIRO (1991): "The Consumption Beta Explains Expected Returns at Long Horizons," Yale University Manuscript.

BREEDEN, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," Journal of Financial Economics, 7, 265–96.

BREEDEN, D. T., M. R. GIBBONS, AND R. H. LITZENBERGER (1989): "Empirical Test of the Consumption-Oriented CAPM," The Journal of Finance, 44(2), 231–262.

CABALLERO, R. J. (1990): "Expenditure on Durable Goods: A Case for Slow Adjustment," Quarterly Journal of Economics, 105, 727–43.

——— (1995): "Near Rationality, Heterogeneity, and Aggregate Consumption," Journal of Money, Credit, and Banking, 27(1), 29–48.

CAMPBELL, J. Y. (1996): "Understanding Risk and Return," Journal of Political Economy, 104(2), 298–345.

——— (1999): "Asset Prices, Consumption, and The Business Cycle," in Handbook of Macroeconomics, ed. by J. B. Taylor, and M. Woodford, vol. 1C, chap. 19, pp. 1231–1303. Elsevier Press.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," Journal of Political Economy, 107(2), 205–51.

Carroll, C. D. (1997): "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," Quarterly Journal of Economics, 112(1), 1–55.

Cochrane, J. H. (1996): "A Cross-Sectional Test of an Investment-Based Asset Pricing Model," Journal of Political Economy, 104(3), 572–621.

Daniel, K. D., and D. Marshall (1997): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle at Long Horizons," Macroeconomic Dynamics, 1(2), 452–84.

Dynan, K. E., and D. M. Maki (2001): "Does the Stock Market Matter for Consumption?," Finance and Economics Discussion Series 2001-23, Federal Reserve Board.

Eichenbaum, M., L. P. Hansen, and K. Singleton (1988): "A Times Series Analysis of Representative Agent Models of Consumption and Leisure," Quarterly Journal of Economics, 103, 51–78.

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

——— (1993): "Common Risk Factors in the Returns on Stocks and Bonds," The Journal of Financial Economics, 33, 3–56.

Fama, E. F., and J. MacBeth (1973): "Risk, Return and Equilibrium: Empirical Tests," Journal of Political Economy, 81, 607–636.

Flavin, M. (1981): "The Adjustment of Consumption to Changing Expectations about Future Income," Journal of Political Economy, 89, 974–1009.

——— (2001): "Owner-Occupied Housing and the Presence of Adjustment Costs: Implications for Asset Pricing and Nondurable Consumption," UCSD working paper.

Gabaix, X., and D. Laibson (2001): "The 6D bias and the equity premium puzzle," in N.B.E.R. Macroeconomics Annual 2001, ed. by B. Bernanke, and K. Rogoff, pp. 257–311. Cambridge: MIT Press.

Gourinchas, P.-O., and J. A. Parker (2002): "Consumption Over the Lifecycle," Econometrica, 70(1), 47–89.

Grossman, S., and G. Laroque (1990): "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," Econometrica, 58(1), 25–51.

Hall, R. E. (1978): "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis," Journal of Political Economy, 86(6), 971–87.

Hall, R. E., and F. S. Mishkin (1982): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, 50, 461–481.

Jagannathan, R., and Z. Wang (1996): "The Conditional CAPM and the Cross-Section of Expected Returns," The Journal of Finance, 51, 3–53.

——— (1998): "An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regression," The Journal of Finance, 53, 1285–1309.

Kan, R. M., and C. Zhang (1999): "Two-pass Tests of Asset Pricing Models with Useless Factors," Journal of Finance, 54, 204–35.

Kandel, S., and R. F. Stambaugh (1990): "Expectations and Volatility of Consumption and Asset Returns," Review of Financial Studies, 3(2), 207–232.

Leamer, E. E. (1983): "Model Choice and Specification," in Handbook of Econometrics, ed. by Z. Grilliches, and M. D. Intriligator, vol. 1. North-Holland.

Lettau, M., and S. Ludvigson (2001a): "Consumption, Aggregate Wealth, and Expected Stock Returns," Journal of Finance, 56(3), 815–49.

——— (2001b): "Resurrectiong the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," Journal of Political Economy, 109, 1238–1286.

Ludvigson, S., and C. Steindel (1999): "How Important is the Stock Market Effect on Consumption?," FRBNY Economic Policy Review, 5(2), 29–51.

Lustig, H. (2001): "The Market Pice of Aggregate Risk and the Wealth Distribution," Stanford University mimeo.

Lynch, A. W. (1996): "Decision Frequency and Synchronization Across Agents: Implications for Aggregate Consumption and Equity Return," Journal of Finance, 51(4), 1479–97.

Mankiw, N. G., and M. D. Shapiro (1986): "Risk and Return: Consumption Beta Versus Market Beta," Review of Economics and Statistics, 68, 452–59.

Marshall, D. A., and N. G. Parekh (1999): "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?," Journal of Finance, 54(2).

Merton, R. (1969): "Lifetime portfolio selection under uncertainty: The continuous time case," Review of Economics and Statistics, 51, 247–57.

Ogaki, M., and C. M. Reinhart (1998): "Measuring Intertemporal Substitution: The Role of Durable Goods," Journal of Political Economy, 106(5), 1078–98.

Parker, J. A. (1999a): "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," American Economic Review, 89(4), 959–973.

——— (1999b): "Spendthrift in America? On two decades of decline in the U.S. saving rate," in N.B.E.R. Macroeconomics Annual, ed. by B. Bernanke, and J. Rotemberg, pp. 317–370. Cambridge: MIT Press.

——— (2001): "The Consumption Risk of the Stock Market," Brookings Papers on Economic Activity, 2, 279–348.

Piazzesi, M. (2001): "Comment on Gabaix and Laibson, 'The 6D bias and the equity premium puzzle'," in N.B.E.R. Macroeconomics Annual 2001, ed. by B. Bernanke, and K. Rogoff. Cambridge: MIT Press.

Piazzesi, M., M. Schneider, and S. Tuzel (2003): "Housing, Consumption, and Asset Pricing," UCLA manuscript.

Rubinstein, M. (1976): "The Valuation of Uncertain Income Streams and the Pricing of Options," Bell Journal of Economics, 7, 407–25.

Shanken, J. (1992): "On the Estimation of Beta-Pricing Models," The Review of Financial Studies, 5, 1–33.

Sims, C. A. (2001): "Implications of Rational Inattention," Princeton University Manuscript.

Souleles, N. (forthcoming): "Consumer Sentiment: Its Rationality and Usefulness in Forecasting Expenditure – Evidence from Michigan Micro Data," Journal of Money, Credit, and Banking.

Startz, R. (1989): "The Stochastic Behavior of Durable and Nondurable Consumption," Review of Economics and Statistics, 71, 356–63.

Wilcox, D. W. (1992): "The Construction of U.S. Consumption Data: Some Facts and Their Implications For Empirical Work," American Economic Review, 82, 922–941.

Zeldes, S. P. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy, 97, 305–346.

Appendixes

A. Multiplicative External Habits

In this section we show that, in the presence of external habit formation that acts as a shifter of the utility function, the ultimate risk to consumption measure is the appropriate stochastic discount factor to price the cross-section of asset returns.

Suppose that the instantaneous marginal utility function, $u'(C_t)$, is shifted by an external habit component so that marginal utility at time t takes the form $V(C_t, \Psi_t) = u'(C_t)\Psi_t(C_0, ..., C_t)$ where Ψ_t is a function of the past history of aggregate consumption. Several properties of the habit component Ψ_t are desirable from an economic point of view: i) Ψ_t is increasing in C_s so that more "habit" consumption increases the marginal utility of agent consumption; ii) the effect of C_{t-s} on Ψ_t decreases with s so that more recent "habit" consumption matters more; iii) the marginal effect of "habit" consumption is decreasing so that C_t has diminishing marginal effects on Ψ_t. We formalize our general class of habit models as follows.

Assumption (The class of habit models)

a) $\Psi_t(C_0, ..., C_{t-\tau}) = f(\ln C_{t-\tau}, \Psi_{t-\tau-1})$ for $\tau \geq 0$

b) $f_C(\ln C, \Psi) \geq 0, \quad f_{CC}(\ln C, \Psi) < 0$

c) $|f_C| > |f_C f_\Psi| \Rightarrow |f_\Psi| < 1$

where f_C and f_Ψ are the derivatives with respect to the first and the second arguments of the function.[19] For example, this class of models includes the model of Abel (1990)

$$u'(C_t)\Psi_t(C_0, ..., C_t) = C_t^{-\gamma}\left(\frac{1}{X_t}\right)^{1-\gamma}$$

where $X_t = C_{t-1}^{\kappa}$ and $1 > \kappa(\gamma - 1) > 0$.

The first order conditions of the agent utility maximization problem imply an arbitrage relation analogous to equations (2.1) and (3.5)

$$E_t\left[u'(C_{t+1+S})\Psi_{t+1+S}R^i_{t,t+1}\right] - E_t\left[u'(C_{t+1+S})\Psi_{t+1+S}\right]R^f = 0. \tag{A.1}$$

Dividing both sides by $u'(C_t)$, taking unconditional expectation and rewriting the expectation of the product in term of covariances and reorganizing yield

$$E\left[R^i_{t,t+1}\right] = \alpha_S - \frac{Cov\left[m_{t+1+S}\Psi_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\Psi_{t+1+S}\right]}, \tag{A.2}$$

[19] This recursive formulation does not represent a substantial loss of generality and greatly reduces the dimensionality of the problem at hand making it easily tractable. The use of the natural logarith is without loss of generality.

where $\alpha_S = R^f$ and $m_{t+1+S} = \frac{u'(C_{t+1+S})}{u'(C_t)}$. Following the same derivation without habit component, yields

$$E\left[R^i_{t,t+1}\right] = \alpha_S - \frac{Cov\left[m_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\right]}. \tag{A.3}$$

Notice that log linearizing the ratio of marginal utilities m_{t+1+S} in equation (A.3) we obtain the beta representation of equation (3.3). This means that, if we are able to show that the right hand side of equation (A.2) converges to the right hand side of equation (A.3), the ultimate risk to consumption is the appropriate stochastic discount factor to price the cross-section of asset returns in the presence of external habits that act as a shifter of the instantaneous utility function.

Proposition 1. *Under the following regularity conditions on consumption growth and the functions*
i) $E\left[\Psi_t\right] < d_0$, $E\left[\Psi_t^2\right] < d_1 \quad \forall t$
ii) $\lim_{t \longrightarrow +\infty} \Psi_t \ln \Psi_t \Longrightarrow \lim_{t \longrightarrow +\infty} E\left[\Psi_t\right] = E\left[\lim_{t \longrightarrow +\infty} \Psi_t\right]$
iii) $\left|\Delta \ln C_{t+1}\right| < d_2 \quad \forall t$ *almost surely*
iv) $\left|f_C\left(C_t, \Psi_{t-1}\right)\right| < d_3 \quad \forall t$ *almost surely*
v) $f_C\left(C, \Psi\right)$ *is* $o\left(t^{-2}\right)$
vi) $Var_t\left(\Delta \ln C_{t+1}\right)$ *is bounded from above for each* t
for $d_0, ..., d_3$ *arbitrary constants,* $\exists\, \overline{S}$ *such that* $\forall S > \overline{S}$

$$\left\| \frac{Cov\left[m_{t+1+S}\Psi_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\Psi_{t+1+S}\right]} - \frac{Cov\left[m_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\right]} \right\| \leqslant 0 \tag{A.4}$$

almost surely and uniformly.

Proof: We prove the proposition for $\tau = 0$. The case for $\tau > 0$ is a simplification. Notice that proving equation (A.4) is equivalent to prove that

$$\frac{E\left[m_{t+1+S}\Psi_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\Psi_{t+1+S}\right]} \overset{a.s.}{\longrightarrow} \frac{E\left[m_{t+1+S}, R^i_{t,t+1}\right]}{E\left[m_{t+1+S}\right]}. \tag{A.5}$$

and the convergence is uniform. We first show that Ψ_{t+1+S} converges almost surely and uniformly as $S \rightarrow +\infty$. By Taylor expansion of $\left|\Psi_{t+1+S} - \Psi_{t+S}\right|$ around Ψ_{t+S-1} and C_{t+S} we obtain that

$$\begin{aligned} \left|\Psi_{t+1+S} - \Psi_{t+S}\right| &= \left|f\left(\ln C_{t+1+S}, \Psi_{t+S}\right) - \Psi_{t+S}\right| \\ &\cong \left|f_C\left(\ln C_{t+S}, \Psi_{t+S-1}\right)\left(\Delta \ln C_{t+S+1}\right) + f_\Psi\left(\ln C_{t+S}, \Psi_{t+S-1}\right)\left(\Psi_{t+S} - \Psi_{t+S-1}\right)\right| \\ &\leq \left|f_C\left(\ln C_{t+S}, \Psi_{t+S-1}\right)\right|\left|\Delta \ln C_{t+S+1}\right| + \left|f_\Psi\left(\ln C_{t+S}, \Psi_{t+S-1}\right)\right|\left|\left(\Psi_{t+S} - \Psi_{t+S-1}\right)\right| \\ &\leq \sum_{\tau=0}^{t+S} \left|f_C\left(\ln C_\tau, \Psi_{\tau-1}\right)\right|\left|f_\Psi\left(\ln C_{\tau+1}, \Psi_\tau\right)\right|^{t+S-\tau}\left|\Delta \ln C_{t+S+1}\right| \qquad (A.6) \\ &\quad + \left|f_\Psi\left(\ln C_1, \Psi_0\right)\right|^{t+S}\left|\left(\Psi_1 - \Psi_0\right)\right| \end{aligned}$$

Notice that the last term in equation (A.6), since $|f_\Psi| < 1$, goes uniformly almost surely to zero as $S \to +\infty$. For any δ_1, δ_2, k, pick T_1 and T_2 such that

$$|f_\Psi|^{t+S} < \delta_1 \frac{\epsilon}{k\,(t+S)} \;\; \forall S \geq T_1$$

$$|f_C\,(C_{t+S}, \Psi_{\tau+S-1})| < \delta_2 \frac{\epsilon}{k\,(t+S)^2} \quad \forall S \geq T_2$$

Then define $\widetilde{T} = \max\,[T_1, T_2]$ and pick $T = 2\widetilde{T}$. It follows that for each ϵ

$$\sum_{\tau=0}^{t+S} |f_C\,(\ln C_\tau, \Psi_{\tau-1})|\,|f_\Psi\,(\ln C_{\tau+1}, \Psi_\tau)|^{t+S-\tau}\,|\Delta \ln C_{t+S+1}| < \epsilon \tag{A.7}$$

almost surely for each $S > T$ since

$$\begin{aligned}
& \sum_{\tau=0}^{t+S} |f_C\,(\ln C_\tau, \Psi_{\tau-1})|\,|f_\Psi\,(\ln C_{\tau+1}, \Psi_\tau)|^{t+S-\tau}\,|\Delta \ln C_{t+S+1}| \\
\leq\;& d_2 \sum_{\tau=t+\widetilde{T}}^{t+S+\widetilde{T}} |f_C\,(\ln C_\tau, \Psi_{\tau-1})|\,|f_\Psi\,(\ln C_{\tau+1}, \Psi_\tau)|^{t+S-\tau}\,|\Delta \ln C_{t+S+1}| \\
& + \sum_{\tau=0}^{t+\widetilde{T}} |f_C\,(\ln C_\tau, \Psi_{\tau-1})|\,|f_\Psi\,(\ln C_{\tau+1}, \Psi_\tau)|^{t+S-\tau}\,|\Delta \ln C_{t+S+1}| \\
<\;& d_2 \left(\delta_1 \frac{\epsilon}{k} + \delta_2 \frac{\epsilon}{k}\right) \frac{1}{(t+S)}
\end{aligned}$$

and this completes the proof of almost sure convergence of Ψ_{t+1+S} Then, taking a log linear approximation of the habit shifter around C_{t+S} and $\Psi_{\tau+S-1}$

$$\begin{aligned}
\Psi_{t+S+1} \;\approx\;& f\,(\ln C_{t+S}, \Psi_{\tau+S-1}) + f_C\,(\ln C_{t+S}, \Psi_{\tau+S-1})\,\Delta \ln C_{t+S+1} \\
& + f_\Psi\,(\ln C_{t+S}, \Psi_{\tau+S-1})\,(\Psi_{\tau+S} - \Psi_{\tau+S-1})
\end{aligned}$$

it is easy to see that, since the conditional on time $t+S$ variance of $\Delta \ln C_{t+S+1}$ is bounded from above and $|f_C\,(C_{t+S}, \Psi_{\tau+S-1})|$ decreases over time due to the increase in consumption,

$$\lim_{S \longrightarrow \infty} Var_{t+S}\,(\Psi_{t+S+1}) = 0$$

almost surely. Therefore, considering a normal approximation of the distribution of Ψ_{t+S+1}, using the approximation $m_{t+1+S} \approx 1 - \gamma \ln \frac{C_{t+S+1}}{C_t}$ and applying the law of iterated expectation in equation (A.5) completes the proof.

Table 1: Expected Returns and Consumption Risk at Different Horizons

Horizon (quarters)	R^2	Î	*Fama-MacBeth standard error*	*Shanken standard error*	*Implied risk free rate (annualized)*	*t-test for the implied risk free rate***	*Implied relative risk aversion*	*95 % confidence interval for implied risk aversion****
0*	0.12	0.0030	(0.0023)	(0.0025)	15.0%	6.08	40.8	[87.2 -35.7]
0	0.19	0.0043	(0.0026)	(0.0030)	8.6%	2.93	55.5	[103.8 -29.6]
1	0.08	0.0037	(0.0039)	(0.0041)	8.7%	2.29	23.1	[53.9 -42.0]
2	0.17	0.0081	(0.0055)	(0.0062)	6.3%	1.41	25.9	[45.9 -23.8]
3	0.20	0.0093	(0.0058)	(0.0065)	6.7%	1.74	20.0	[34.0 -12.9]
4	0.11	0.0073	(0.0064)	(0.0067)	8.0%	2.05	12.9	[25.6 -19.2]
5	0.12	0.0085	(0.0072)	(0.0077)	8.0%	2.11	11.9	[22.6 -18.3]
6	0.11	0.0091	(0.0078)	(0.0083)	8.1%	2.12	10.7	[20.0 -17.9]
7	0.16	0.0116	(0.0082)	(0.0089)	7.2%	1.86	11.0	[18.6 -12.0]
8	0.28	0.0165	(0.0084)	(0.0097)	6.4%	1.74	12.4	[18.2 -4.1]
9	0.41	0.0213	(0.0084)	(0.0103)	6.2%	1.68	13.0	[17.5 0.8]
10	0.44	0.0234	(0.0087)	(0.0110)	5.8%	1.51	12.6	[16.4 1.5]
11	**0.52**	**0.0276**	**(0.0090)**	**(0.0118)**	**5.3%**	**1.27**	**12.5**	**[15.7 3.5]**
12	0.42	0.0299	(0.0108)	(0.0146)	4.8%	1.06	12.0	[15.1 0.7]
13	0.45	0.0392	(0.0128)	(0.0195)	3.6%	0.62	12.4	[15.0 0.3]
14	0.38	0.0409	(0.0140)	(0.0215)	4.4%	0.89	11.7	[14.3 -1.9]
15	0.37	0.0479	(0.0157)	(0.0261)	4.1%	0.74	11.6	[13.9 -4.4]

* Uses the start of period timing convention: consumption growth between t+1 and t+2 is aligned with returns during t+1.

** Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the return on 3 month Treasury bills.

*** The confidence interval for risk aversion uses the Shanken standard error.

Table 2. Pricing Errors of Average Size and Book to Market Portfolios

By size quintiles
(Quarterly rates in percentage terms)

	Horizon (quarters) 0	11	*Change*
S1	0.87	0.72	-0.15
S2	0.60	0.40	-0.21
S3	0.42	0.23	-0.18
S4	0.35	0.29	-0.05
S5	0.18	0.20	0.02

By book-to-market quintiles
(Quarterly rates in percentage terms)

	Horizon (quarters) 0	11	*Change*
B1	0.87	0.70	-0.17
B2	0.40	0.25	-0.14
B3	0.27	0.35	0.08
B4	0.35	0.23	-0.12
B5	0.58	0.36	-0.22

Note: Average pricing errors are calculated as the square root of the average squared errors. Units are quarterly rates reported in percentage terms, so that an improvement of 0.15 is roughly 0.6 percent per year

Table 3: Expected Returns and Consumption Risk Estimated 1963Q3 to 1998Q3

Horizon (quarters)	R^2	Î	*Fama-MacBeth standard error*	*Shanken standard error*	*Implied risk free rate (annualized)*	*t-test for the implied risk free rate***	*Implied relative risk aversion*	*95 % confidence interval for implied risk aversion****
0*	0.03	0.0008	(0.0021)	(0.0021)	12.3%	3.87	12.3	[61.4 -65.6]
0	0.34	0.0048	(0.0025)	(0.0029)	6.3%	1.37	60.7	[106.5 -18.2]
1	0.14	0.0054	(0.0048)	(0.0052)	5.2%	0.78	30.9	[63.7 -51.3]
2	0.29	0.0125	(0.0074)	(0.0095)	2.2%	0.00	34.9	[55.9 -41.1]
3	0.26	0.0133	(0.0085)	(0.0104)	4.1%	0.55	25.4	[41.3 -33.7]
4	0.14	0.0095	(0.0087)	(0.0095)	6.1%	1.13	15.7	[30.0 -37.2]
5	0.15	0.0110	(0.0097)	(0.0107)	6.0%	1.12	14.3	[26.1 -34.8]
6	0.23	0.0160	(0.0110)	(0.0130)	4.6%	0.70	15.6	[24.7 -27.7]
7	0.26	0.0178	(0.0114)	(0.0136)	3.8%	0.45	14.4	[22.0 -21.7]
8	0.45	0.0272	(0.0121)	(0.0166)	2.5%	0.07	16.1	[21.6 -10.6]
9	0.64	0.0352	(0.0115)	(0.0177)	2.1%	0.04	16.4	[20.4 -0.3]
10	0.65	0.0376	(0.0122)	(0.0190)	0.9%	0.35	15.4	[18.9 -0.5]
11	**0.73**	**0.0426**	**(0.0119)**	**(0.0196)**	**0.9%**	**0.34**	**14.8**	**[17.7 3.0]**
12	0.70	0.0477	(0.0137)	(0.0239)	-1.1%	0.78	14.2	[16.8 0.0]
13	0.70	0.0563	(0.0148)	(0.0284)	-1.6%	0.83	13.9	[16.1 -0.5]

* Uses the start of period timing convention: consumption growth between t+1 and t+2 is aligned with returns during t+1.

** Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the return on 3 month Treasury bills.

*** The confidence interval for risk aversion uses the Shanken standard error.

Table 4: Expected Returns and the Risk to Total Consumption

Horizon (quarters)	R^2	$\hat{l}$	*Fama-MacBeth standard error*	*Shanken standard error*	*Implied risk free rate (annualized)*	*t-test for the implied risk free rate***	*Implied relative risk aversion*	*95 % confidence interval for implied risk aversion****	
0*	0.27	0.0053	(0.0026)	(0.0031)	14.6%	5.67	49.5	[80.7	-12.0]
0	0.32	0.0053	(0.0024)	(0.0027)	9.1%	3.37	49.8	[78.4	-2.2]
1	0.23	0.0064	(0.0037)	(0.0041)	7.0%	2.26	27.7	[45.1	-12.7]
2	0.18	0.0083	(0.0054)	(0.0060)	6.3%	1.69	19.6	[32.3	-16.5]
3	0.17	0.0103	(0.0069)	(0.0077)	6.6%	1.87	15.9	[25.7	-16.6]
4	0.11	0.0085	(0.0073)	(0.0078)	7.7%	2.17	10.9	[19.7	-20.9]
5	0.16	0.0133	(0.0090)	(0.0101)	6.9%	1.98	11.9	[18.5	-16.2]
6	0.17	0.0140	(0.0092)	(0.0103)	7.1%	2.16	10.6	[16.1	-13.0]
7	0.20	0.0177	(0.0103)	(0.0120)	7.0%	2.23	10.5	[15.0	-10.7]
8	0.33	0.0236	(0.0104)	(0.0129)	6.6%	2.16	10.9	[14.1	-2.7]
9	0.41	0.0271	(0.0103)	(0.0132)	6.3%	2.01	10.4	[13.0	0.6]
10	0.46	0.0306	(0.0108)	(0.0143)	6.3%	1.98	9.9	[12.1	1.6]
11	**0.55**	**0.0370**	**(0.0111)**	**(0.0157)**	**6.2%**	**1.82**	**9.7**	**[11.4**	**3.4]**
12	0.50	0.0411	(0.0126)	(0.0185)	6.1%	1.77	9.3	[10.8	2.4]
13	0.44	0.0520	(0.0158)	(0.0264)	5.4%	1.47	9.2	[10.6	-0.5]
14	0.38	0.0531	(0.0169)	(0.0280)	6.3%	1.82	8.7	[10.0	-2.1]
15	0.37	0.0588	(0.0182)	(0.0317)	7.0%	2.00	8.4	[9.5	-3.8]

* Uses the start of period timing convention: consumption growth between t+1 and t+2 is aligned with returns during t+1.

** Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the return on 3 month Treasury bills.

*** The confidence interval for risk aversion uses the Shanken standard error.

Table 5: Expected Returns and Consumption Risk Using Equal Weighted Returns

Horizon (quarters)	R^2	$\hat{I}$	*Fama-MacBeth standard error*	*Shanken standard error*	*Implied risk free rate (annualized)*	*t-test for the implied risk free rate***	*Implied relative risk aversion*	*95 % confidence interval for implied risk aversion****
0*	0.10	0.0031	(0.0023)	(0.0025)	15.9%	6.12	42.6	[88.5 -32.9]
0	0.39	0.0064	(0.0028)	(0.0036)	7.0%	2.04	74.7	[122.2 -14.2]
1	0.23	0.0069	(0.0041)	(0.0047)	5.7%	1.18	37.2	[64.8 -22.1]
2	0.38	0.0131	(0.0058)	(0.0076)	2.7%	0.27	35.7	[53.4 -10.7]
3	0.39	0.0141	(0.0063)	(0.0078)	3.9%	0.68	26.3	[39.0 -5.0]
4	0.30	0.0130	(0.0068)	(0.0080)	4.6%	0.83	19.5	[30.2 -7.8]
5	0.31	0.0153	(0.0078)	(0.0093)	4.6%	0.82	17.7	[26.6 -7.4]
6	0.31	0.0166	(0.0086)	(0.0102)	4.5%	0.79	15.9	[23.5 -7.5]
7	0.36	0.0200	(0.0092)	(0.0114)	3.5%	0.52	15.4	[21.6 -4.6]
8	0.50	0.0242	(0.0092)	(0.0121)	3.6%	0.62	15.3	[20.2 0.1]
9	0.62	0.0288	(0.0091)	(0.0127)	4.0%	0.78	15.0	[19.0 3.5]
10	0.67	0.0326	(0.0097)	(0.0141)	3.2%	0.49	14.6	[17.9 4.1]
11	**0.74**	**0.0362**	**(0.0097)**	**(0.0147)**	**3.0%**	**0.43**	**14.0**	**[16.8 5.3]**
12	0.67	0.0414	(0.0120)	(0.0192)	1.8%	0.05	13.5	[16.2 2.5]
13	0.71	0.0545	(0.0144)	(0.0271)	0.2%	-0.41	13.8	[16.0 0.2]
14	0.67	0.0589	(0.0159)	(0.0310)	0.9%	-0.22	13.1	[15.2 -3.0]
15	0.68	0.0689	(0.0176)	(0.0380)	0.5%	-0.33	12.8	[14.7 -8.7]

* Uses the start of period timing convention: consumption growth between t+1 and t+2 is aligned with returns during t+1.

** Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the return on 3 month Treasury bills.

*** The confidence interval for risk aversion uses the Shanken standard error.

Table 6: The Significance of Consumption Risk vs. Firm Characteristics

Horizon (quarters)	R^2	$\hat{l}$	*Coefficient on Size Factor*	R^2	$\hat{l}$	*Coefficient on Book to Market Factor*
			Factor calculated as mean of the log			
0	0.22	0.0068	0.0012	0.56	0.0028	0.0048
		(0.0025)	(0.0013)		(0.0029)	(0.0019)
		(0.0033)	(0.0018)		(0.0030)	(0.0020)
11	0.58	0.0355	0.0012	0.65	0.0186	0.0033
		(0.0106)	(0.0014)		(0.0115)	(0.0022)
		(0.0158)	(0.0021)		(0.0133)	(0.0025)
			Factor calculated as log of the mean			
0	0.23	0.0071	0.0012	0.78	0.0013	0.0079
		(0.0026)	(0.0013)		(0.0030)	(0.0025)
		(0.0034)	(0.0017)		(0.0030)	(0.0026)
11	0.59	0.0357	0.0012	0.78	0.0075	0.0071
		(0.0106)	(0.0013)		(0.0140)	(0.0034)
		(0.0158)	(0.0020)		(0.0144)	(0.0035)

Note: Standard errors in parenthesis, the first based on the Fama-McBeth standard error while the second uses the Shanken correction.

Table 7: Statistics on Predicting Average Returns for Betas from Different Factor Models

Sample period of asset returns: 1963:3-1998:3

			Factors							*t-test for the implied risk free rate**	*Implied relative risk aversion***
Row	R^2	R^f	R^m_{t+1}	SMB_{t+1}	HML_{t+1}	cay_t	$\Delta \ln C_{t+1}$	$cay_t \Delta \ln C_{t+1}$	$\ln(C_{t+1+s}/C_t)$		
1	0.78	3.9%	0.0118	0.4425	1.4666						
		(6.3)	(0.0165)	(0.5099)	(0.4477)					0.2627	
		(6.7)	(0.0173)	(0.5108)	(0.4477)					0.2485	
2	0.66	13.6%				-0.0032	0.0035	0.0022			
		(3.8)				(0.0035)	(0.0020)	(0.0012)		3.1239	84.6
		(5.1)				(0.0046)	(0.0027)	(0.0016)		2.3087	[5258.8 -5258.8]
3	0.43	8.27%					0.0019	0.0012			
		(2.8)					(0.0025)	(0.0015)		2.1780	49.7
		(3.6)					(0.0032)	(0.0020)		1.6890	[5365.5 -5266.3]
4	0.73	0.9%							0.0426		
		(4.1)							(0.0119)	0.3354	14.8
		(6.9)							(0.0196)	0.1998	[3.0 17.7]
5	0.34	6.3%					0.0048				
		(3.0)					(0.0025)			1.3691	
		(3.5)					(0.0029)			1.1742	

Sample period of asset returns: 1952:4-1998:3

			Factors							*t-test for the implied risk free rate**	*Implied relative risk aversion***
Row	R^2	R^f	R^m_{t+1}	SMB_{t+1}	HML_{t+1}	cay_t	$\Delta \ln C_{t+1}$	$cay_t \Delta \ln C_{t+1}$	$\ln(C_{t+1+s}/C_t)$		
1	0.71	16.0%	-0.0116	0.3266	1.2820						
		(5.4)	(0.0130)	(0.4070)	(0.3751)					2.7258	
		(5.5)	(0.0134)	(0.4075)	(0.3748)					2.6375	
2	0.61	15.7%				-0.0065	0.0049	0.0030			
		(4.3)				(0.0037)	(0.0019)	(0.0011)		3.3491	65.1
		(6.0)				(0.0051)	(0.0026)	(0.0016)		2.3878	[5894.7 -5764.5]
3	0.47	9.4%					0.0018	0.0012			
		(2.6)					(0.0027)	(0.0016)		2.9520	28.0
		(3.3)					(0.0035)	(0.0021)		2.2821	[7347.5 -7291.6]
4	0.61	5.1%							0.0306		
		(3.3)							(0.0091)	0.9415	13.1
		(4.6)							(0.0125)	0.6757	[4.8 16.1]
5	0.33	7.5%					0.0044				
		(2.7)					(0.0020)			2.0432	
		(3.1)					(0.0023)			1.7955	

Note: Standard errors in parenthesis, the first based on the Fama-McBeth standard error while the second uses the Shanken correction.

* Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the 3 months T-bill. The top statistic uses Fama-MacBeth standard errors, the bottom statistic uses the Shanken correction.

** 95% confidence interval in square brackets (confidence interval for risk aversion uses the Shanken standard error).

Table 8: Explaining Expected Returns with Consumption Growth and the Risk-Free Interest Rate

Horizon (quarters)	R^2 *-- two factors model*	R^2 *-- one factor model*	$\hat{l}^C$	*Fama-MacBeth standard error*	*Shanken standard error*	*Implied risk free rate (annualized)*	*t-test for the implied risk free rate***	*Implied relative risk aversion*	*95 % confidence interval for implied risk aversion****	
0*	0.12	0.12	0.0030	(0.0023)	(0.0025)	15.0%	6.08	40.8	[87.2	-35.7]
0	0.19	0.19	0.0043	(0.0026)	(0.0030)	8.6%	2.93	55.5	[103.8	-29.6]
1	0.15	0.08	0.0038	(0.0039)	(0.0041)	8.6%	2.27	23.4	[53.8	-41.0]
2	0.24	0.18	0.0083	(0.0055)	(0.0062)	6.1%	1.36	26.1	[45.7	-22.3]
3	0.27	0.21	0.0096	(0.0058)	(0.0065)	6.5%	1.68	20.1	[33.7	-11.6]
4	0.18	0.12	0.0076	(0.0064)	(0.0068)	7.8%	1.99	13.0	[25.3	-18.1]
5	0.19	0.13	0.0090	(0.0074)	(0.0079)	7.8%	2.04	12.0	[22.3	-16.8]
6	0.18	0.12	0.0098	(0.0080)	(0.0085)	7.8%	2.03	10.8	[19.7	-16.0]
7	0.22	0.17	0.0125	(0.0084)	(0.0092)	7.0%	1.78	11.0	[18.2	-10.5]
8	0.32	0.30	0.0177	(0.0087)	(0.0101)	6.3%	1.68	12.2	[17.6	-3.2]
9	0.44	0.43	0.0229	(0.0087)	(0.0108)	6.1%	1.64	12.6	[16.8	1.3]
10	0.47	0.45	0.0253	(0.0092)	(0.0116)	5.7%	1.47	12.1	[15.7	1.9]
11	**0.55**	**0.54**	**0.0299**	**(0.0095)**	**(0.0126)**	**5.3%**	**1.27**	**11.9**	**[14.8**	**3.6]**
12	**0.46**	**0.44**	**0.0328**	**(0.0115)**	**(0.0156)**	**4.7%**	**1.04**	**11.3**	**[14.2**	**1.2]**
13	0.50	0.47	0.0427	(0.0134)	(0.0206)	3.7%	0.66	11.6	[14.0	1.1]
14	0.43	0.41	0.0457	(0.0150)	(0.0233)	4.4%	0.87	11.0	[13.3	-0.6]
15	0.42	0.41	0.0536	(0.0168)	(0.0283)	4.0%	0.72	10.8	[12.8	-2.3]

* Uses the start of period timing convention: consumption growth between t+1 and t+2 is aligned with returns during t+1.

** Tests for the null hypothesis that the implied risk free rate is equal to the time series average of the return on 3 month Treasury bills.

***The confidence interval for risk aversion uses the Shanken standard error.

Figure 1: Ratio of Statistical Variance to Cross-sectional Variance of Beta's

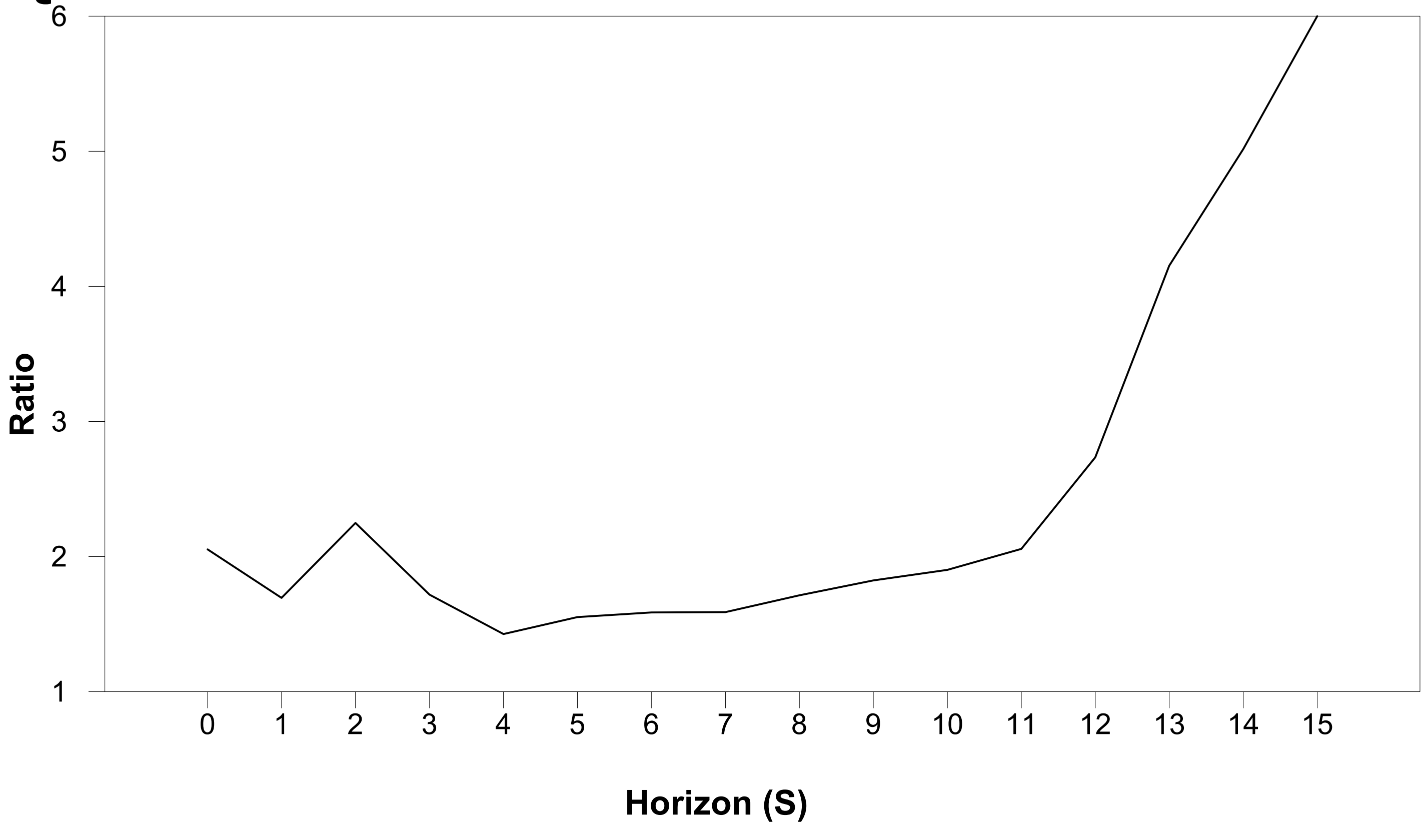

Figure 2: Average Returns and Betas

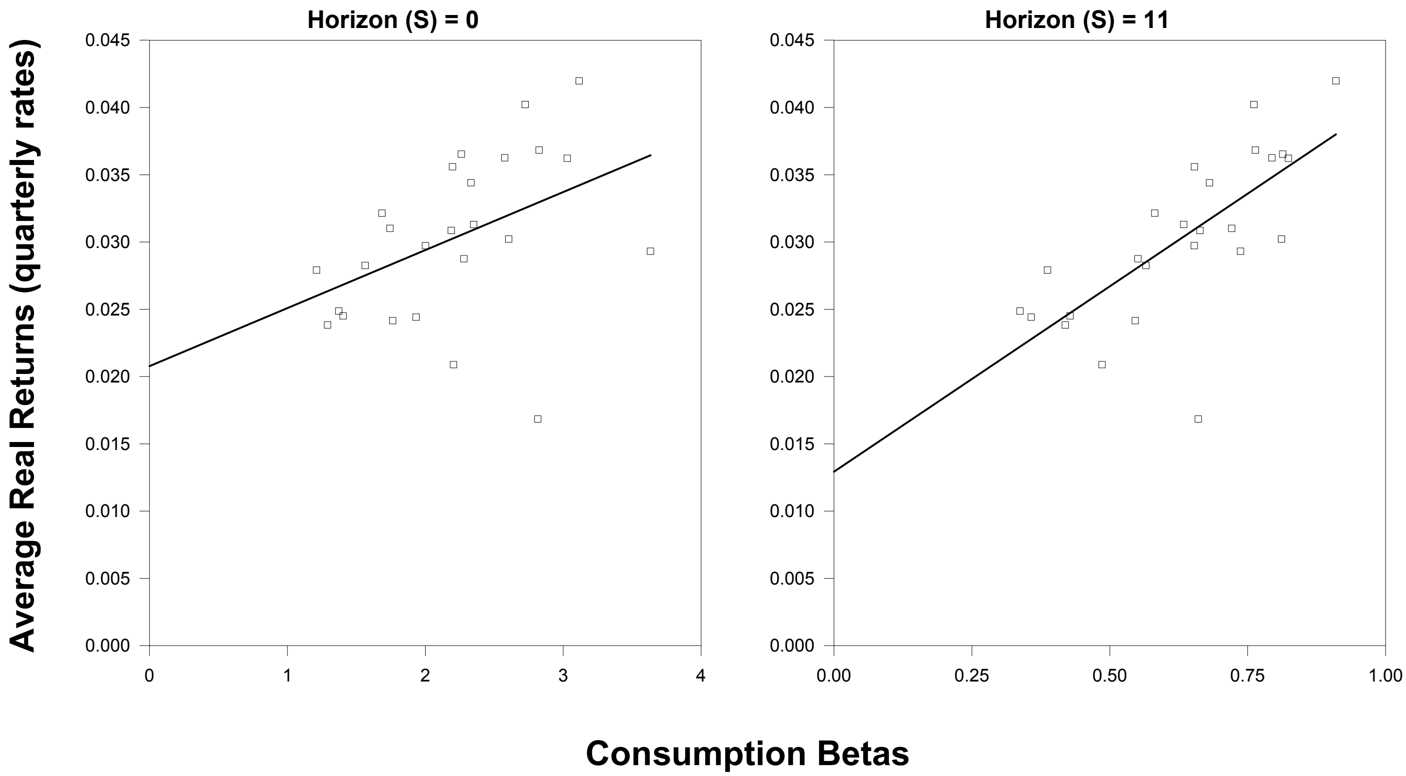

Figure 3: Fitted Returns and Average Returns

Quarterly rates

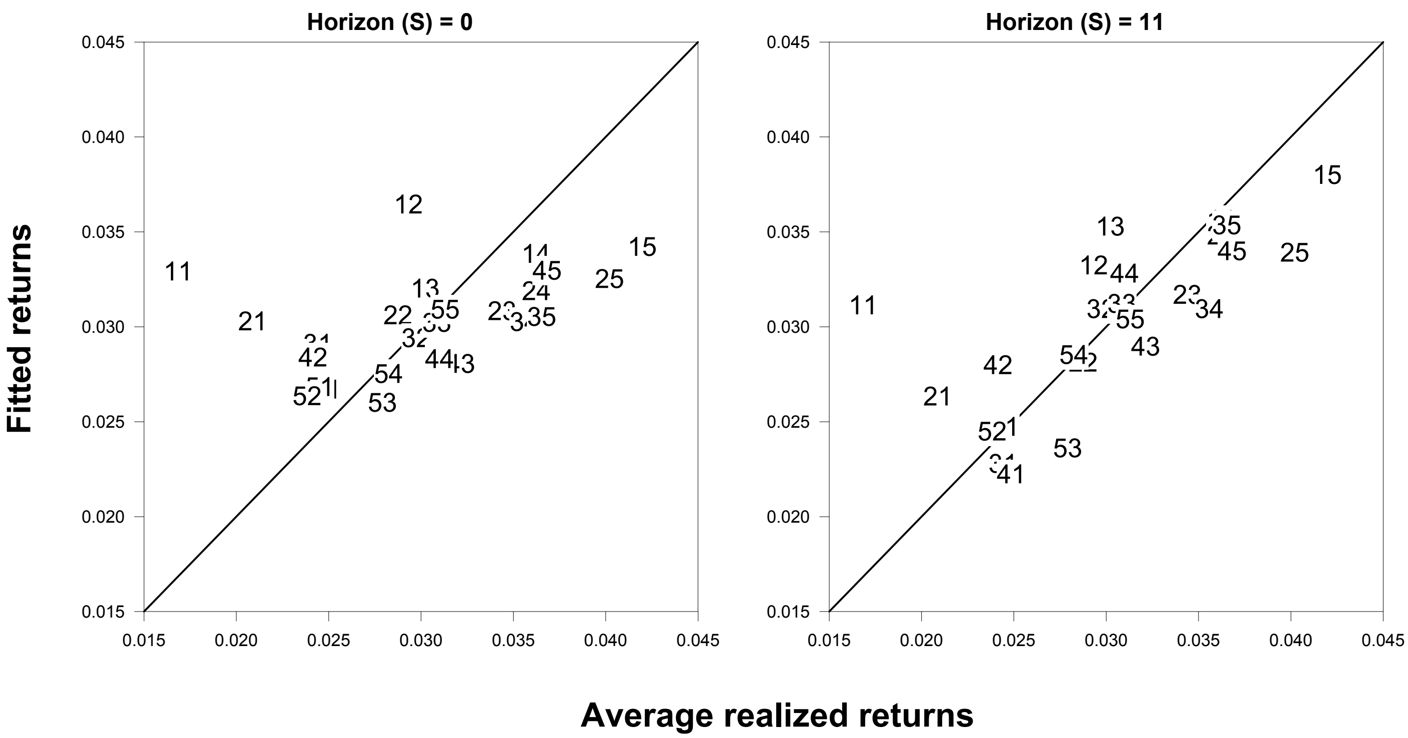

Figure 4: Fitted Returns and Average Returns for Different Models

Quarterly rates, 1963:3 - 1998:3

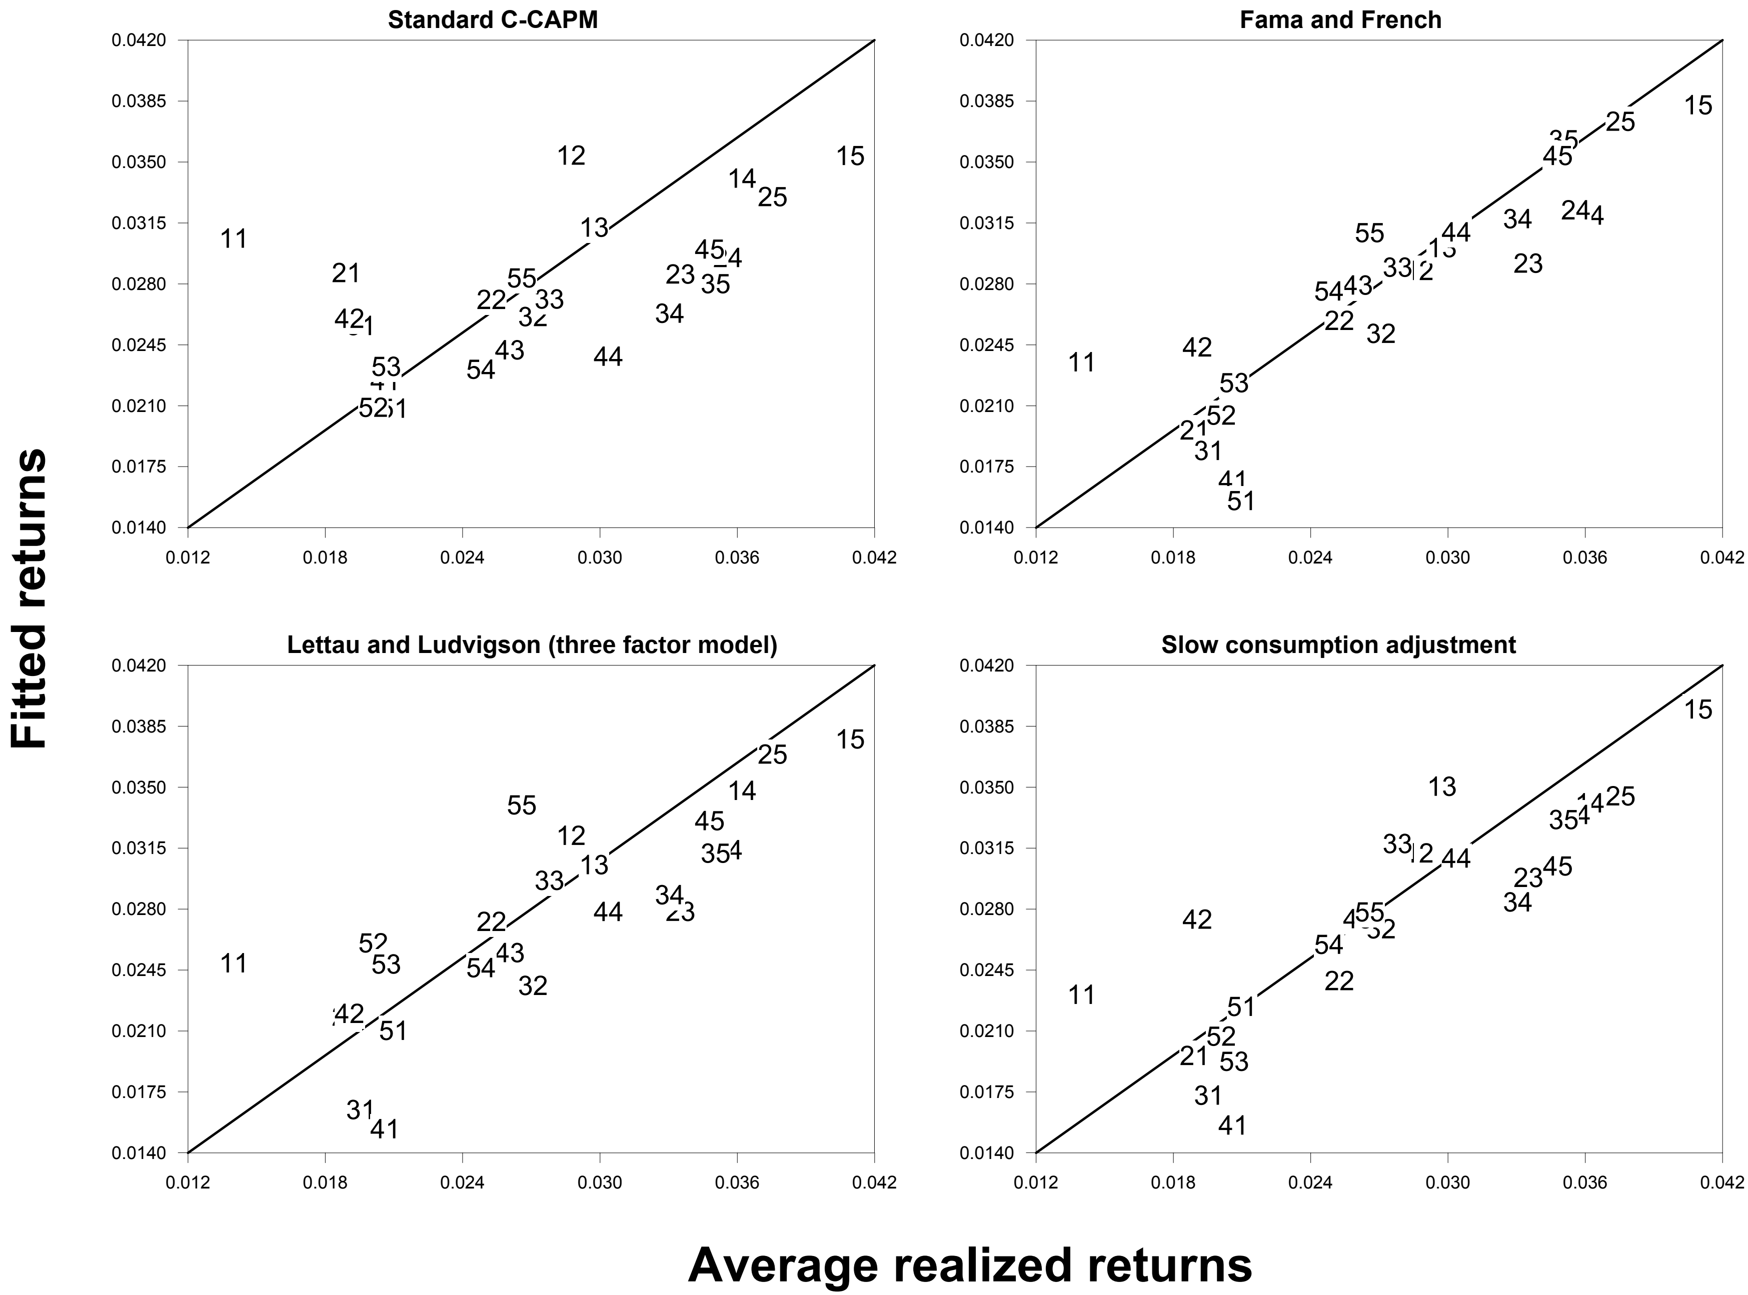

Figure 5: Fitted Returns and Average Returns in the Baseline Sample

Quarterly rates, 1952:4 - 1998:3