

Consumption Risk and the Cross Section of Expected Returns

Jonathan A. Parker

Princeton University and National Bureau of Economic Research

Christian Julliard

Princeton University

This paper evaluates the central insight of the consumption capital asset pricing model that an asset's expected return is determined by its equilibrium risk to consumption. Rather than measure risk by the contemporaneous covariance of an asset's return and consumption growth, we measure risk by the covariance of an asset's return and consumption growth cumulated over many quarters following the return. While contemporaneous consumption risk explains little of the variation in average returns across the 25 Fama-French portfolios, our measure of ultimate consumption risk at a horizon of three years explains a large fraction of this variation.

I. Introduction

The natural economic explanation for differences in expected returns across assets is differences in risk. According to canonical economic theory, the risk of an asset is determined by its covariance with con-

For helpful comments and discussions, we thank Yacine Ait-Sahalia, John Campbell, John Cochrane, Kent Daniel, Albina Danilova, Pierre-Olivier Gourinchas, Sydney Ludvigson, Monika Piazzesi, Christopher Sims, Motohiro Yogo, two anonymous referees, and seminar participants at the University of Chicago, Princeton University, and the NBER Asset Pricing workshop, July 2003. For financial support, Parker thanks the National Science Foundation (grant SES-0096076) and the Sloan Foundation. We thank Kenneth French for making the Fama-French portfolio data available on his Web page: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[*Journal of Political Economy*, 2005, vol. 113, no. 1]
© 2005 by The University of Chicago. All rights reserved. 0022-3808/2005/11301-0007\$10.00

sumption growth. But differences in the covariance of returns and contemporaneous consumption growth across portfolios do not explain the differences in expected returns observed in the U.S. stock market (Mankiw and Shapiro 1986; Breeden, Gibbons, and Litzenberger 1989; Campbell 1996; Cochrane 1996; Lettau and Ludvigson 2001*b*). The asset pricing literature has largely concluded that differences in expected returns are not due to differences in risk to consumption, but instead 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 reevaluate the central insight of the consumption capital asset pricing model (CCAPM) that an asset's expected return is determined by its equilibrium risk to consumption. Rather than measure the risk of a portfolio by the contemporaneous covariance of its return and consumption growth—as done in the previous literature on the CCAPM and the cross-sectional pattern of expected returns—we 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 risk on the basis of 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 has a long history of relatively successful application in economics. Second, this approach is consistent with the canonical CCAPM in that, if the CCAPM were true, the ultimate risk would correctly measure the risk of a portfolio. Finally and most importantly, the ultimate risk may be a better measure of the true risk of an asset if consumption is slow to adjust to returns. If consumption responds with a lag to changes in wealth, then the contemporaneous covariance of consumption and wealth understates or mismeasures the true risk of a portfolio. Ultimate consumption risk, on the other hand, can provide the correct measure of risk under several extant explanations of slow consumption adjustment, such as some models of (*a*) measurement error in consumption; (*b*) costs of adjusting consumption; (*c*) nonseparability of the marginal utility of consumption from factors such as labor supply or housing stock, which themselves are constrained to adjust slowly; or (*d*) constraints on information flow or calculation so that household behavior is “near-rational.” 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 expected returns even though the true stochastic discount factor—the complete model of household saving and portfolio choice—has to date escaped discovery.

Our main finding is that ultimate consumption risk can largely explain the cross-sectional pattern of expected portfolio returns. While the covariance of each portfolio and *contemporaneous* consumption growth explains little of the variation in expected returns across portfolios, at a horizon of three years the *ultimate* risk to consumption explains from 44 to 73 percent of the variation in expected returns across portfolios, depending on specification. The performance of ultimate consumption risk as a linear one-factor model rivals that of the Fama and French (1993) three-factor model and the Lettau and Ludvigson (2001*b*) three-factor model, two important extant linear factor models that have been used to price the expected returns in the Fama-French portfolios. Despite the fact that we focus on a horizon at which consumption growth performs well, our results are unlikely to be due to chance. There are statistically significant differences across portfolios in their covariation with consumption growth. And in a Monte Carlo experiment, repeated estimation on data sets with no serial correlation in consumption growth rarely finds that ultimate consumption risk prices the portfolios much better than contemporaneous consumption risk.

What drives our results? The success of ultimate consumption risk comes from the fact that the excess returns on the Fama-French portfolios predict future consumption growth. The Fama-French portfolios mainly capture the size premium and the value premium, and both the excess return of small firms less large firms and the excess return of high-value stocks less low-value stocks predict consumption growth. The joint significance of these two excess returns in predicting consumption growth peaks at the horizon of three years, where the fit of the ultimate consumption model of expected returns is largest.

The serial correlation of consumption growth gives the ultimate consumption risk stochastic discount factor a clear business cycle pattern. Consumption falls through recessions, so that the stochastic discount factor is highest right before and at the start of recessions. Value stocks, and to a lesser extent small stocks, have high average returns because they pay off poorly before and early in recessions, and ultimate consumption risk captures this.

Despite the strong relationship between ultimate consumption risk and expected returns, ultimate consumption risk does not provide a complete accounting of relative expected returns. The model has 26 moments and three parameters, and generalized method of moments (GMM) tests of overidentification reject the contemporaneous and ultimate consumption risk models. We note that consumption risk prices poorly the expected return on the portfolio with the smallest size and value, suggesting that issues of liquidity or the inability to short may be important elements missed by consumption risk alone.

We build on the large literature testing consumption-based asset pric-

ing models, but in particular, our main results are most closely related to those of Brainard, Nelson, and Shapiro (1991), which shows that the longer the horizon of the investor, the better the CCAPM performs relative to the CAPM, and to those of Bansal, Dittmar, and Lundblad (forthcoming), which shows 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 estimate structural parameters. More recently, Lustig and Van Nieuwerburgh (2003), Piazzesi, Schneider, and Tuzel (2003), and Yogo (2003) price expected returns using durable consumption or housing consumption in conjunction with nondurable consumption.¹ Finally, our article is closely related to the literature on the stochastic properties of aggregate consumption following aggregate market returns (Kandel and Stambaugh 1990; Daniel and Marshall 1997; Ludvigson and Steindel 1999; Parker 2000, 2001; Dynan and Maki 2001; Gabaix and Laibson 2002; Piazzesi 2002; Bansal and Yaron 2004).

The balance of the paper is organized as follows. Section II uses the optimality condition for the portfolio choice of a representative agent to derive the ultimate consumption risk version of the CCAPM, then briefly discusses alternative models for which ultimate consumption risk and not contemporaneous risk would determine expected returns. Section III describes our data and Section IV our econometric methodology. Section V contains our main results on fit, significance, and implied risk aversion for the ultimate risk to consumption. Section VI lays out what lies behind the main findings—the predictability of consumption growth and the business cycle pattern of our stochastic discount factor—and presents the results of a simulation exercise that demonstrates that our findings are unlikely to be due merely to chance. Section VII of the paper compares a one-factor linear model of ultimate consumption risk to the fit and the performance of the models of Fama and French (1993) and Lettau and Ludvigson (2001*b*). Section VIII presents conclusions.

II. Expected Returns and Consumption Risk

The CCAPM, first developed by Rubinstein (1976) and Breeden (1979), relates expected returns to their equilibrium risk to consumption. A representative household allocates its resources among consumption and different investment opportunities so as to maximize the expected present discounted value of utility flows from consumption. When

¹ We suspect that these models work well because expenditure levels on durable goods and the prices of durable goods react more quickly to news than nondurable consumption, and so signal or predict future movements in nondurable consumption.

wealth is allocated optimally across assets, the marginal investment in any assets yields the same expected increase in future utility, so that for any excess return $R_{i,t+1}^e$,

$$E_t[u'(C_{t+1})R_{i,t+1}^e] = 0, \quad (1)$$

where $u(\cdot)$ is the period utility function, C_{t+1} is consumption, and i indexes excess returns.

We can write this equation as a model of expected returns by dividing by $u'(C_t)$ and using the unconditional expectations operator and the definition of covariance to yield

$$E[R_{i,t+1}^e] = -\frac{\text{Cov}[m_{t+1}, R_{i,t+1}^e]}{E[m_{t+1}]}, \quad (2)$$

where $m_{t+1} = u'(C_{t+1})/u'(C_t)$ is the stochastic discount factor. Relative expected excess returns on different portfolios are determined by their relative covariances with the stochastic discount factor and, thus, by their relative comovement with consumption. A portfolio that has greater consumption risk in equilibrium has greater expected return, since consumption and marginal utility are inversely related.

According to equations (1) and (2), differences in risk across portfolios are due to differences in their *contemporaneous* comovement with consumption. This maintains several assumptions not directly related to portfolio choice: the agent must perfectly smooth expected marginal utility over time, marginal utility must be determined only by consumption, and the level of consumption must be costless to adjust. Under these conditions, the impact of any return is reflected instantly and completely in consumption.

Instead of maintaining these assumptions, we evaluate the central insight of the CCAPM—that consumption risk determines average returns—while allowing for the possibility that consumption is slow to respond to returns.

Consider the following alternative model of expected returns, derived under the assumption that the CCAPM is literally true. Use the consumption Euler equation for the risk-free rate between $t+1$ and $t+1+S$,

$$u'(C_{t+1}) = E_{t+1}[\delta R_{t+1,t+1+S}^f u'(C_{t+1+S})], \quad (3)$$

to eliminate $u'(C_{t+1})$ from equation (1) to yield

$$E_t[m_{t+1}^S R_{i,t+1}^e] = 0, \quad (4)$$

where $m_{t+1}^S = R_{t+1,t+1+S}'u'(C_{t+1+S})/u'(C_t)$. When we reorganize as before, expected returns are given by

$$E[R_{i,t+1}^e] = -\frac{\text{Cov}[m_{t+1}^S, R_{i,t+1}^e]}{E[m_{t+1}^S]}, \quad (5)$$

where we refer to $-\text{Cov}[m_{t+1}^S, R_{i,t+1}^e]$ for large S as *ultimate* consumption risk. A portfolio's risk is measured by the covariance of its return at $t+1$ and the change in marginal utility from t to $t+1+S$, where S is the horizon over which the consumption response is studied.

Why measure risk and price expected returns using ultimate consumption risk, as in equation (5), instead of contemporaneous consumption risk, as in equation (2)? 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 the long-term measure provides a better accounting of the equilibrium risk of different stocks than the contemporaneous measure does. In general, there are three classes of reasons why consumption might be slow to respond: measurement error in consumption, mismeasurement of marginal utility due to nonseparabilities with other factors, and slow adjustment of marginal utility itself.

First, 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 real aggregate consumption from monthly survey data on nominal sales at retail establishments. Second, the marginal utility of consumption may be altered by fluctuations in such factors as hours of leisure, habits, housing stock, durable consumption, and so forth.² If these factors are stationary and covary with returns, their impact can be similar to that of transitory measurement error, confounding inference using equation (2) whereas ultimate consumption risk still determines returns. Third, consumption may be slow to adjust to wealth shocks because of the presence of constraints on information flow or direct costs of adjusting consumption.³

² An example that drops expected utility theory is given by Restoy and Weil (1998), which shows that under Kreps-Porteus-Epstein-Zin preferences, expected returns are priced by the covariance of the return with contemporaneous consumption growth plus a factor proportional to the covariance with revisions to future consumption growth.

³ For examples of the second type of model, see Eichenbaum, Hansen, and Singleton (1988), Startz (1989), Abel (1990), and Flavin (2001). For examples of the third type of model, see Caballero (1995), Lynch (1996), Marshall and Parekh (1999), Alvarez, Atkeson, and Kehoe (2002), and Gabaix and Laibson (2002). We provide more detail on these arguments in Parker and Julliard (2003).

III. Data

We use the quarterly returns on the 25 Fama and French (1992) portfolios $\{R_{i,t+1}\}_{i=1}^{25}$ and construct excess returns as these returns less the return on a three-month Treasury bill, $R_{t,t+1}^f$. We study these returns because the Fama-French portfolios have a large dispersion in average returns that is relatively stable in subsamples and because they have been used extensively to evaluate asset pricing models. These portfolios are designed to focus on two features of average returns: the size effect—firms with small market value have, on average, higher returns—and the value premium—firms with high book values relative to market equity have, on average, higher returns.

More specifically, the 25 Fama-French portfolios are the intersections of five portfolios formed on size (market equity) and five portfolios formed on the ratio of book equity to market equity. Data on portfolio returns are available monthly from July 1926 to December 2002. We denote a portfolio by the rank of its market equity and then the rank of its book-to-market ratio so that portfolio 15 is the smallest quintile of stocks by market equity and the largest quintile of stocks by book-to-market. To match the frequency of consumption data, we cumulate returns to a quarterly frequency, so that $R_{i,t+1}$ represents the return on portfolio i during the quarter $t + 1$.

For consumption, we use real (chain-weighted) personal consumption expenditures on nondurable goods per capita from the National Income and Product Accounts (NIPA). Except where noted, we make the standard “end-of-period” timing assumption that consumption during quarter t takes place at the end of the quarter, so that $\text{Cov}[m_{t+1}^0, R_{i,t+1}^e]$ is calculated using NIPA consumption in $t + 1$ relative to t and returns during $t + 1$. We make this choice mostly because under this convention the entire period that C_t covers is contained in the information set of the agent before $R_{i,t+1}^e$, so that as we increase S we omit no covariance of consumption and returns. The alternative timing convention, used by Campbell (1999), for example, is that consumption occurs at the beginning of the period, so that, using NIPA dates, one aligns m_{t+2}^0 with $R_{i,t+1}^e$. For comparison, for $S = 0$, we also report results for this alternative. All returns are deflated by the same deflator as consumption.

The sample of returns that we use is limited by the availability of consumption data to run from the second quarter of 1947 to the fourth quarter of 1999. The sample of returns stops at the end of 1999, so that we can allow up to four 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, we use all available (not preliminary) consumption data, from the first quarter of 1947 to the fourth quarter of 1999 plus S quarters, which is the third quarter of 2003 when $S = 15$.

IV. Estimation Methodology

We estimate the model for different horizons (S) by choosing the parameters to make the pattern of expected excess returns based on consumption risk as close as possible to the observed pattern of average returns. We parameterize the stochastic discount factor by assuming that the utility function exhibits constant relative risk aversion with risk aversion parameter γ_S ,

$$m_t^S = R_{t,t+S}^f \left(\frac{C_{t+S}}{C_{t-1}} \right)^{-\gamma_S},$$

and estimate the parameters μ_S , γ_S , and α_S for each S by GMM using the 26×1 empirical moment function

$$g(\mathbf{R}_t^e, C_{t+1+S}, C_{t-1}; \mu_S, \gamma_S, \alpha_S) = \begin{bmatrix} \mathbf{R}_t^e - \alpha_S \mathbf{1}_{25} + \frac{(m_t^S - \mu_S) \mathbf{R}_t^e}{\mu_S} \\ m_t^S - \mu_S \end{bmatrix}, \quad (6)$$

where $\mathbf{R}_t^e$ is the 25×1 vector whose i th element is $R_{i,t}^e$. Equation (5) implies that the moment function satisfies the 26 moment restrictions

$$E[g(\mathbf{R}_t^e, C_{t+1+S}, C_{t-1}; \mu_S, \gamma_S, \alpha_S)] = 0 \quad (7)$$

at the true parameter values. For inference, we rely on the asymptotic distributions for GMM, assuming that the time dimension is becoming infinite and the maximum S is fixed.⁴ It is worth making two points about this approach to estimation.

First, we base estimation on equation (5) rather than equation (4) even though this choice means having to include the additional moment $E[m_t^S - \mu_S] = 0$ because choosing these moments allows different models (different S) to be evaluated using a similar criterion. For any stochastic discount factor, the differences between the empirical and theoretical moments are pricing errors: the extent to which the expected return predicted by the model does not equal the observed average excess return.⁵ Thus the units of these errors are independent of the choice of stochastic discount factor.

Second, by including the parameter α_S rather than imposing $\alpha_S = 0$, we separately evaluate the ability of the model to explain the equity premium and the cross section of expected stock returns. Because of

⁴ We construct the 26×26 covariance matrix of the moments using the Newey and West (1987) procedure with $S+1$ cross correlations in the t dimension, to be conservative about statistical uncertainty. We find very similar results either fixing the cross correlations at four for all horizons or using the VARHAC procedure.

⁵ More precisely, these are errors in expected return. However, since they are all scaled by the mean of the stochastic discount factor, they are proportional to pricing errors, and we maintain this terminology.

our choice of moments, α_s is in units of expected return and so measures the extent to which the model underpredicts the excess returns of all Fama-French portfolios by the same amount. That is, it measures the extent to which the model has an equity premium puzzle. If we omitted α_s , then we might incorrectly conclude that ultimate consumption risk was only weakly related to expected returns across portfolios when in fact ultimate consumption risk was “merely” not consistent with the average excess return of all portfolios. In fact, many potential explanations of the equity premium—such as limited participation, differential taxation of stocks and bonds, liquidity demand for Treasury bills, and changing regulation of asset markets—can be consistent with consumption risk pricing the expected returns among stocks, but not between stocks and Treasury bills.⁶

We report estimates both from GMM with a prespecified weighting matrix and from efficient GMM (iterated to convergence). The prespecified weighting matrix is a diagonal matrix that places weight one on the first 25 moments and very large weight on the last moment.⁷ This estimator has three advantages over efficient GMM. First, given this weighting matrix, these estimates match the mean of the stochastic discount factor and minimize the sum of squared pricing errors on the Fama-French portfolios, giving each portfolio equal weight. Thus this choice of weighting matrix forces the model to try to explain the size effect and the value premium. Efficient GMM, on the other hand, minimizes the sum of squared pricing errors on weighted combinations of the portfolios, focusing on linear combinations of returns that have low variance and often ignoring the value premium or size effect or both if they are “hard” to price. In practice, efficient GMM prices rather unusual combinations of portfolios, with extreme long and short positions. Second, because GMM with a prespecified weighting matrix tries to price the same portfolios as one varies S , measures of fit and specification tests are more comparable across different models (different S) than for efficient GMM (Cochrane 2001, chap. 11). Third, GMM with a prespecified weighting matrix has superior small-sample properties (see, e.g., Ferson and Foerster 1994; Hansen, Heaton, and Yaron 1996; Ahn and Gadarowski 1999).

For our estimates using the prespecified weighting matrix, we follow Jagannathan and Wang (1996) and Hansen and Jagannathan (1997)

⁶ We checked that the inclusion of α does not drive our results by defining excess returns relative to portfolio 33. The parameter restriction $\alpha = 0$ is not rejected, and we reach the same main conclusions, with the exception that we are unable to evaluate the consistency of the fitted model with the equity premium.

⁷ The weight of the last moment is chosen large enough that significant variation in the weight does not change the parameter estimates. This ensures that our findings are not due in any way to misestimating the mean of the stochastic discount factor as we increase S . Yogo (2003) shows the importance of including the last moment.

and evaluate the models' performance using the Hansen-Jagannathan distance and its p -value.⁸ The Appendix contains the details of this test statistic.

V. Consumption Risk and Expected Returns

This section asks whether consumption risk explains the cross-sectional variation in expected returns on different portfolios of stocks. First, is consumption risk economically significant—does consumption risk explain a large share of the variance of average returns? Second, is consumption risk statistically significant? Third, is consumption risk a complete explanation of expected returns—do tests of overidentification reject the model?

Our estimates provide two additional pieces of information about the model. First, we estimate the risk aversion of the representative investor. This is a structural parameter and should be consistent with behavior under risk in other economic environments. Second, our estimates of α_s measure the extent to which the relative consumption risk of different portfolios is consistent with the average excess return on all portfolios.

We begin by estimating the model using the prespecified weighting matrix. The first row of table 1 presents the results for the contemporaneous CCAPM (horizon $S = 0$) using the “beginning-of-period” timing convention for consumption growth that aligns $(C_{t+2}/C_{t+1})^{-\gamma}$ with $R_{i,t+1}^e$, and the second row reports results using the “end-of-period” timing convention that aligns $(C_{t+1}/C_t)^{-\gamma}$ with $R_{i,t+1}^e$. The contemporaneous CCAPM performs poorly in four ways.

First, contemporaneous consumption risk is not an economically significant determinant of the cross section of expected returns. Column 1 displays the percentage of the variation in average returns explained by the fitted model, given by the cross-sectional R^2 .⁹ Consumption risk

⁸ We present the Hansen-Jagannathan distance rather than the first-stage J -test because we find the former to be more numerically stable (the latter involves a difficult matrix pseudo-inversion). As seen in the tables, the Hansen-Jagannathan test and the efficient J -test lead to similar inferences.

⁹ This is a “standard” R^2 calculated as though we had done a nonlinear regression of the consumption covariances on the average returns:

$$R^2 = 1 - \frac{\text{Var}(E_T[R_i^e] - \hat{R}_i^e)}{\text{Var}(E_T[R_i^e])},$$

where

$$E_T[x_i] := \frac{1}{T} \sum_{t=1}^T x_i$$

and

$$\hat{R}_i^e = \hat{\alpha}_s + \frac{\text{Cov}_T[\hat{m}_{t+1}^s, R_{i,t+1}^e]}{\hat{m}^s} = \hat{\alpha}_s + \frac{E_T[(\hat{m}_i^s - \hat{\mu}_s) R_{i,t}^e]}{\hat{\mu}_s}.$$

TABLE 1
EXPECTED EXCESS RETURNS AND DIFFERENT HORIZONS OF CONSUMPTION RISK

HORIZON S (Quarters)	GMM WITH PRESPECIFIED WEIGHTING MATRIX				EFFICIENT GMM			
	R^2 (%) (1)	α (2)	Risk Aversion		R^2 (%) (5)	α (6)	Risk Aversion	
			γ (3)	Dist (4)			γ (7)	J Test $\chi^2(23)$ (8)
0*	4	.029 (.006)	19.9 (33.3)	.37 [.000]	−32	.024 (.005)	−47.9 (19.0)	46.7 [.002]
0	3	.023 (.005)	19.0 (41.8)	.37 [.000]	3	.024 (.005)	17.0 (21.7)	52.2 [.000]
1	2	.023 (.007)	10.7 (27.5)	.37 [.000]	1	.024 (.006)	2.3 (16.9)	52.1 [.000]
2	5	.020 (.009)	14.6 (24.8)	.37 [.000]	3	.023 (.006)	4.6 (14.8)	52.2 [.000]
3	10	.018 (.009)	17.9 (23.5)	.36 [.000]	3	.023 (.006)	3.5 (13.6)	51.9 [.001]
4	4	.021 (.008)	9.1 (17.2)	.37 [.000]	3	.023 (.007)	6.2 (12.1)	51.3 [.001]
5	7	.019 (.008)	11.7 (16.3)	.36 [.000]	6	.022 (.006)	8.9 (10.3)	50.4 [.001]
6	9	.018 (.008)	12.6 (15.3)	.36 [.000]	2	.024 (.006)	2.3 (9.5)	51.9 [.001]
7	10	.019 (.008)	11.0 (14.3)	.36 [.000]	5	.023 (.006)	3.5 (8.7)	51.7 [.001]
8	20	.018 (.006)	15.1 (13.8)	.34 [.000]	10	.023 (.006)	5.2 (8.6)	51.5 [.001]
9	30	.018 (.005)	17.9 (12.5)	.31 [.000]	21	.023 (.006)	8.7 (7.8)	51.6 [.001]
10	33	.017 (.005)	18.6 (13.7)	.31 [.000]	24	.022 (.006)	9.3 (7.9)	51.4 [.001]
11	44	.015 (.006)	25.4 (16.4)	.28 [.000]	38	.020 (.006)	15.8 (8.6)	49.3 [.001]
12	32	.016 (.005)	25.0 (16.5)	.31 [.000]	15	.022 (.006)	7.3 (7.9)	52.4 [.000]
13	35	.012 (.006)	38.5 (14.0)	.30 [.000]	14	.022 (.006)	8.0 (8.7)	52.3 [.000]
14	30	.014 (.005)	34.6 (24.6)	.31 [.000]	7	.023 (.006)	5.0 (8.0)	52.4 [.000]
15	24	.016 (.008)	39.4 (24.4)	.33 [.000]	1	.024 (.006)	1.7 (8.7)	52.7 [.000]

NOTE.—GMM estimation of eqq. (5) and (6). Standard errors are reported in parentheses and p -values in brackets. GMM with a prespecified weighting matrix uses an identity matrix except that the weight on the last moment is large, as discussed in Sec. IV. Efficient GMM iterates until convergence. Covariance matrices are calculated using the Newey-West procedure with $S+1$ lags.

* Uses the beginning-of-period timing convention. Denotes using the timing convention for consumption data that consumption takes place at the start of each quarter.

explains only 3 or 4 percent of the cross-sectional variation in average returns. Second, given the estimated levels of contemporaneous consumption risk, the average return is far too large. The estimated intercept is statistically significant and implies that the average excess return on a Fama-French portfolio exceeds that implied by its contemporaneous consumption risk by roughly 9 or 12 percent per year.¹⁰

Third, the model is rejected by the data. Column 4 presents the Hansen-Jagannathan distance and the probability that one would err in rejecting the model on the basis of this distance, that is, the p -value of a specification test based on this distance. The Hansen-Jagannathan distance would be the square root of a weighted average of the squared pricing errors if we did not include a moment for the mean of m_{t+1}^s . Since we do, this interpretation of the Hansen-Jagannathan distance as a measure of average pricing error is not strictly correct; however, in this case this interpretation is not misleading. The expected return based on the fitted model is off by (roughly) 0.37 percent per quarter for the “typical” portfolio.

Despite the low fit of the model, contemporaneous consumption risk is statistically significant in explaining expected returns across portfolios. Given the differences in the contemporaneous consumption risk across portfolios, the point estimate for the risk aversion of the representative agent required to rationalize the spread in average returns is near 20. A Wald-type test treating the GMM estimator as an extremum estimator rejects the restriction that $\gamma = 0$ at very high levels of significance (p -value $< 10^{-15}$; not shown in table 1).¹¹ The standard errors on risk aversion are so large because the GMM objective function is quite flat in γ near the estimate and for larger γ but becomes increasingly steep as $\gamma \rightarrow 0$.

The remaining rows that display GMM estimates with the prespecified weighting matrix show that, contrary to the result for contemporaneous risk, consumption risk measured after consumption has had time to adjust to returns explains a significant share of the variance in average returns. As the horizon over which consumption adjustment is measured increases, the economic importance of consumption risk rises up to $S = 11$, where ultimate consumption risk explains 44 percent of the variance in average returns across portfolios. This is in fact the smallest peak explanatory power across specifications that we explore: ultimate consumption risk explains one-half to three-quarters of the variation for alternative specifications (tables 3 and 5 below). The explanatory power of consumption risk is lower for $S \geq 12$, a feature discussed subsequently.

¹⁰ This is consistent with the well-documented poor performance of the CCAPM in explaining the excess return on the market (Grossman and Shiller 1981; Hansen and Singleton 1982; Mehra and Prescott 1985).

¹¹ The test follows Hayashi (2000, 489–90; eqq. 7.4.11, 7.4.9).

For now, we analyze the explanatory power of ultimate consumption risk at three years ($S = 11$). Figure 1 plots the predicted and average returns of different portfolios for the contemporaneous and ultimate consumption risk measures. The horizontal distance between a portfolio and the 45-degree line is the extent to which the expected return based on fitted consumption risk (on the horizontal axis) differs from the observed average return (on the vertical axis). For $S = 0$, there is almost no relation between predicted and average returns. For $S = 11$, the model fits most portfolios quite well, with the exception of the small firms (denoted 11, 12, ..., 15), particularly the smallest size and book-to-market portfolio (11), which should have a much greater expected return based on its consumption risk. For all but six of the 25 portfolios (12, 13, 32, 33, 43, and 53), the pricing error is smaller for $S = 11$ than for $S = 0$.

Even though small, low-value firms are priced poorly, there is still substantial improvement in nearly all value and size categories in moving from horizon $S = 0$ to $S = 11$. Table 2 shows the square root of the average squared pricing error for each size and book-to-market quintile at each horizon. The pricing errors are smaller at $S = 11$ for every group of portfolios except the middle book-to-market portfolios (book-to-market 3). For these portfolios, the pricing error associated with contemporaneous consumption risk is the smallest of the portfolio groups, and the observed increase in the pricing error from $S = 0$ to $S = 11$ is smaller in magnitude than all but one of the decreases in pricing errors in table 2. For small firms, which have the largest pricing errors at both horizons, predicted expected returns are 0.5 percent per year closer to average returns for ultimate consumption risk compared to contemporaneous risk.

In addition to a greater economic role for ultimate consumption risk relative to contemporaneous risk, the statistical significance of m_i^S rises as the horizon S increases to 11 (again, not shown in table 1). While $\gamma_S = 0$ is always strongly rejected, the χ^2 statistics testing this null hypothesis increase by more than two orders of magnitude from $S = 0$ to $S = 11$.

Despite the economic significance of ultimate consumption risk, the data reject that this single factor is the only determinant of expected returns. Table 1 shows that the p -values on the distance tests reject the model for all S . However, as can be inferred from comparing the distance statistics across rows, the p -values, while extremely low, rise up to $S = 11$. Figure 1 shows that portfolio 11 is the most poorly priced, suggesting that the missing element in the CCAPM may be an account of the costs of short selling or the thinness of the market (see D'Avolio 2002; Lamont and Thaler 2003).

It is also worth noting that, as we increase the horizon, there is an




FIG. 1.—Fitted and average returns of consumption risk models: *a*, horizon $S = 0$; *b*, horizon $S = 11$. All returns are quarterly rates. Each portfolio is denoted by the rank of its market equity and then the rank of its ratio of book value to market value. Fitted values are based on the model estimates of table 1 with a prespecified weighting matrix.

TABLE 2
PRICING ERRORS OF AVERAGE SIZE AND VALUE PORTFOLIOS

	HORIZON S (Quarters)		CHANGE
	0	11	
Market equity quintiles:			
Market equity 1	.78	.68	-.11
Market equity 2	.58	.38	-.20
Market equity 3	.37	.22	-.15
Market equity 4	.37	.28	-.10
Market equity 5	.28	.16	-.12
Book-to-market quintiles:			
Book-to-market 1	.74	.65	-.09
Book-to-market 2	.35	.23	-.11
Book-to-market 3	.17	.27	.10
Book-to-market 4	.35	.19	-.15
Book-to-market 5	.70	.41	-.28

NOTE.—Based on estimates of table 1 with the prespecified weighting matrix. Units are quarterly rates reported in percentage terms. Average pricing errors are calculated as the square root of the average squared errors.

improvement in the extent to which the model underpredicts all excess returns. While the intercept remains statistically significant—so that there remains an equity premium puzzle—the magnitude of the puzzle is reduced.¹² Finally, as S rises initially, estimated risk aversion declines, although it rises again to levels around 20 when the fit of the model is greatest (horizons around three years) and becomes greater still for longer horizons.

When we reweight the portfolios (moments) efficiently, ultimate consumption risk still performs best at $S = 11$ in terms of its economic significance in explaining the expected returns on the original portfolios. Estimated risk aversion is significantly lower (more plausible) at all horizons and is estimated more precisely. Otherwise, our conclusions using efficient GMM estimates remain quite similar to those using first-stage GMM.

The balance of the paper demonstrates both the robustness and the source of our main finding that ultimate consumption risk does a better job than contemporaneous consumption risk of pricing expected returns. Section VI explains what lies behind the success of ultimate consumption risk, shows that this result is unlikely to be due only to chance, and addresses whether there is a “best” horizon S . Section VII linearizes our model and compares its performance to two existing successful linear factor pricing models. Before turning to these results, however,

¹² There is already significant evidence that large S does not “solve” the equity premium puzzle in aggregate consumption data (Kandel and Stambaugh 1990; Daniel and Marshall 1997; Parker 2001; Gabaix and Laibson 2002; Piazzesi 2002).

we show that not only is our main result robust to some variations in data, but it also is typically stronger in alternative specifications.

We consider four alternative ways of estimating the importance of consumption risk. In each case, ultimate consumption risk explains more of the cross section of expected returns than in our baseline specification. In each case, the pattern of findings with horizon is similar to that of table 1, so we limit our reporting to contemporaneous consumption risk and horizons near $S = 11$.

First, much previous work has focused on a shorter time period than we analyze in our baseline results.¹³ Panel A of table 3 shows the results of estimating our model on a sample of returns that starts in the third quarter of 1963, a starting period set to match that of Lettau and Ludvigson (2001*b*). In this subperiod, the pattern of coefficients and fit tell a similar story, except that ultimate consumption risk does even better at explaining expected returns. Around 60 percent of the variation in expected returns is explained by consumption risk over a horizon around three years at a level of risk aversion around 30. And the fitted model overstates the average return on all portfolios by less than in the baseline sample.

Second, we measure consumption risk using total consumption instead of nondurable consumption. Ait-Sahalia, Parker, and Yogo (2004) argue that the consumption risk of equity is understated by NIPA nondurable goods because it contains many necessities and few luxury goods. 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.

Panel B of table 3 shows that using total consumption risk in place of nondurable consumption risk leads to broadly similar conclusions. Ultimate consumption risk using total consumption fits the cross section of expected returns slightly better than nondurable consumption, finding slightly higher risk aversion in the first-stage GMM estimates and slightly lower risk aversion using the efficient procedure.

Third, we consider a slightly different set of returns: the 25 equal-weighted Fama-French portfolios. As shown in panel C, ultimate consumption risk does an even better job of explaining the cross-sectional pattern of expected returns of these (albeit similar) portfolios. Not only is over two-thirds of the variation in expected returns explained, but

¹³ The Fama-French portfolios had, and no longer have, the shortcoming that the sample of firms changed significantly starting in 1963 because of a limited availability of the book value of common equity prior to 1962.

TABLE 3
THE ROBUST RELATIONSHIP BETWEEN EXPECTED RETURNS AND CONSUMPTION RISK

HORIZON <i>S</i> (Quarters)	GMM WITH PRESPECIFIED WEIGHTING MATRIX				EFFICIENT GMM			
	<i>R</i> ² (%) (1)	α (2)	Risk Aversion	Dist (4)	<i>R</i> ² (%) (5)	α (6)	Risk Aversion	<i>J</i> -Test $\chi^2(23)$ (8)
			γ (3)				γ (7)	
A. Original Fama-French Start Date: 1963:3–2003:3								
0*	0	.023 (.006)	−2.0 (42.1)	.36 [.00]	0	.020 (.005)	.7 (24.0)	49.9 [.001]
0	21	.017 (.009)	70.5 (50.7)	.32 [.00]	16	.020 (.005)	87.3 (22.2)	48.9 [.001]
10	51	.011 (.006)	21.0 (15.1)	.25 [.000]	23	.019 (.006)	6.9 (7.3)	51.3 [.001]
11	61	.010 (.006)	27.4 (19.1)	.23 [.000]	42	.017 (.005)	13.0 (7.8)	53.1 [.000]
12	63	.009 (.006)	33.8 (21.8)	.22 [.000]	37	.016 (.005)	13.2 (8.1)	54.7 [.000]
B. Total Consumption								
0*	10	.030 (.006)	28.8 (29.1)	.36 [.00]	−5	.024 (.005)	−7.2 (17.0)	51.8 [.001]
0	13	.024 (.007)	26.4 (28.9)	.35 [.00]	8	.023 (.005)	11.9 (16.8)	51.7 [.001]
10	55	.015 (.007)	32.4 (12.8)	.25 [.000]	29	.022 (.005)	9.1 (6.9)	50.6 [.001]
11	58	.019 (.006)	31.2 (21.5)	.24 [.000]	34	.021 (.006)	10.5 (7.2)	51.0 [.001]
12	39	.020 (.006)	25.7 (14.0)	.29 [.000]	21	.021 (.006)	7.4 (6.5)	52.9 [.000]
C. Fama-French Equal-Weighted Portfolios								
0*	4	.030 (.006)	20.1 (30.8)	.45 [.00]	−66	.025 (.005)	−48.2 (18.2)	55.4 [.000]
0	29	.025 (.008)	42.1 (25.7)	.38 [.00]	27	.026 (.005)	34.3 (14.2)	59.2 [.000]
10	73	.013 (.011)	42.5 (23.2)	.24 [.000]	62	.022 (.005)	25.1 (9.7)	45.8 [.003]
11	67	.020 (.007)	34.0 (24.4)	.26 [.000]	65	.021 (.005)	29.3 (10.5)	44.5 [.005]
12	52	.020 (.007)	31.3 (21.9)	.32 [.000]	53	.018 (.005)	29.2 (10.7)	46.4 [.003]
D. Long-Horizon Returns								
0*	4	.029 (.006)	19.9 (33.3)	.37 [.00]	−32	.024 (.005)	−47.9 (19.0)	46.7 [.002]
0	3	.023 (.005)	19.0 (41.8)	.37 [.00]	3	.024 (.005)	17.0 (21.7)	52.2 [.000]
10	49	.121 (.121)	69.2 (61.2)	4.19 [.000]	42	.029 (.009)	100.4 (21.8)	195.7 [.000]
11	52	.158 (.099)	90.2 (56.5)	4.55 [.000]	48	.077 (.003)	162.5 (33.1)	210.0 [.000]
12	43	.261 (.074)	55.4 (31.6)	5.60 [.000]	−362	.102 (.000)	392.8 (24.3)	211.0 [.000]

NOTE.—See the note to table 1.

* Uses the beginning-of-period timing convention.

estimated risk aversion and fit are more stable between first-stage and efficient GMM.

Fourth, we use ultimate consumption risk to price long-horizon returns. That is, instead of substituting the consumption Euler equation for the risk-free rate into the usual optimality condition for portfolio choice, we substitute the consumption Euler equation for the excess return, giving the condition

$$E\left[\left(\frac{C_{t+1+S}}{C_t}\right)^{-\gamma_S} \prod_{\tau=t+1}^{t+1+S} R_{i,\tau}^e\right] = 0.$$

Ex ante, we expect this model to perform worse, since to the extent that consumption responds slowly to returns (or events correlated with returns), this model does not include the consumption response to events that occur near $t + 1 + S$. Panel D of table 3 shows that in fact we find a similar fit for some horizons for this model. Pricing long-horizon expected returns, we find much higher levels of risk aversion than in our baseline model, and the efficient GMM estimates are less consistent with the first-stage estimates and highly unstable across horizons. The model is also very strongly rejected. That said, the model does quite well in terms of fit.

Having shown that the ultimate risk to consumption explains much of the cross-sectional pattern of expected returns, we now ask how and why.

VI. What Drives This Result?

This section demonstrates that ultimate consumption risk does a better job of pricing the portfolios than contemporaneous risk because the returns on the Fama-French portfolios predict future consumption growth. The risk-free interest rate in the discount factor plays almost no role. Consumption is close to a martingale. But because the contemporaneous covariance between returns and consumption growth is so small, a small amount of predictability of consumption growth, in the right pattern across assets, leads to a large increase in the relationship between consumption risk and expected returns with S .

Moreover, these results are unlikely to be due to chance. There are statistically significant differences across portfolios in their covariation with consumption growth. We show that, as the horizon S increases, the difference in the covariance with consumption of small returns relative to big returns grows, both statistically and economically; and similarly for high-value relative to low-value returns. And in a Monte Carlo experiment, repeated estimation on data sets with no serial correlation in

consumption growth rarely finds that ultimate consumption risk prices the portfolios much better than contemporaneous consumption risk.

The serial correlation of consumption growth gives the ultimate consumption risk stochastic discount factor a clear business cycle pattern. Consumption falls through recessions, so that the stochastic discount factor is highest right before and at the start of recessions. Thus value stocks and, to a lesser extent, small stocks have high average returns because they pay off poorly before recessions.

Finally, we take up the question of whether there is a correct or best horizon of consumption adjustment S to measure ultimate consumption risk. On the basis of the theoretical motivation for our measure, it is reasonable to select the “best horizon” by increasing S until one no longer adds more signal than noise to the covariances of returns and consumption growth. Arguably, such a criterion leads one to select $S = 11$.

To demonstrate these claims, log-linearize m_{t+1}^S and assume the approximation that the risk-free rate is constant and equal to the discount rate ($\delta R^f = 1$ in eq. [3]), so that equation (5) can be written in terms of consumption growth:¹⁴

$$\begin{aligned} E[R_{i,t+1}^e] &\approx \gamma \text{Cov} \left[\ln \left(\frac{C_{t+1+s}}{C_t} \right), R_{i,t+1}^e \right] \\ &= \gamma \sum_{s=0}^S \text{Cov} [\Delta \ln (C_{t+1+s}), R_{i,t+1}^e]. \end{aligned} \quad (8)$$

Any expected excess return depends on the sum of covariances of the return with current and future consumption growth rates. If consumption growth were a martingale, then contemporaneous consumption risk and ultimate consumption risk would perform equivalently at pricing the Fama-French portfolios.

We focus on the size and value premia rather than trying to describe and infer the relationship between all 25 excess returns and consumption growth. We use the two excess returns constructed by Fama and French (1993) that capture these premia: the excess return on a portfolio containing stocks of firms with high ratios of book value to market equity relative to a portfolio of firms with low book value to market equity (“high minus low,” denoted HML), and the excess return on a portfolio containing stocks of small firms relative to a portfolio of large firms (“small minus big,” denoted SMB). These are two of the three factors in the Fama-French three-factor model studied in Section VII.

¹⁴ The assumption that the risk-free real interest rate equals the discount rate omits only a small factor, $\{1 - \gamma E[\ln (C_{t+1+s}/C_t)]\}^{-1}$. Evaluated at the estimated $\hat{\gamma}_S$, this factor equals 1.04 at $S = 0$ and 1.15 at $S = 11$.

TABLE 4
PREDICTABILITY OF CONSUMPTION GROWTH BY SIZE AND VALUE RETURNS

HORIZON S (Quarters)	REGRESSION					JOINT SIGNIFICANCE $\chi^2(2)$ (6)	VARIANCE OF $100 \times$ $\ln(C_{t+1+s}/C_t)$ (7)
	$\ln(C_{t+1+s}/C_t)$ on HML_{t+1} and SMB_{t+1}	R^2 (%)	HML $_{t+1}$ on $\ln(C_{t+1+s}/C_{t+1})$	R^2 (%)	SMB $_{t+1}$ on $\ln(C_{t+1+s}/C_{t+1})$		
	R^2 (%) (1)	R^2 (%) (2)	Coefficient (3)	R^2 (%) (4)	Coefficient (5)		
0	1.78						.64
1	2.72	.14	-.259 (.522)	1.11	.705 (.464)	2.6 [.279]	1.35
2	2.06	.00	-.018 (.284)	.84	.422 (.342)	1.5 [.467]	2.22
3	2.51	.07	.094 (.226)	1.35	.418 (.280)	2.4 [.301]	3.22
4	3.39	.01	-.029 (.185)	2.23	.446 (.231)	3.7 [.153]	4.12
5	3.02	.00	-.017 (.152)	1.91	.364 (.213)	2.9 [.230]	4.96
6	2.84	.01	.023 (.150)	1.75	.319 (.193)	2.8 [.251]	5.69
7	2.95	.05	.052 (.139)	1.83	.304 (.174)	3.2 [.203]	6.45
8	2.47	.39	.135 (.128)	1.10	.221 (.163)	3.0 [.227]	6.94
9	2.44	.77	.183 (.119)	.81	.184 (.150)	3.9 [.144]	7.45
10	2.56	.93	.194 (.112)	.79	.175 (.141)	4.6 [.103]	7.84
11	2.51	1.47	.238 (.107)	.43	.127 (.134)	5.8 [.054]	8.39
12	1.23	.49	.132 (.103)	.27	.095 (.126)	2.2 [.329]	9.03
13	.55	.17	.076 (.101)	.07	.047 (.118)	.7 [.694]	9.63
14	.39	.12	.061 (.095)	.04	.033 (.115)	.5 [.779]	10.32
15	.15	.09	.051 (.094)	.01	-.012 (.113)	.3 [.859]	10.86

NOTE.—Col. 1 reports the fit of a regression of contemporaneous and future consumption growth for different S on both returns; cols. 2 and 3 report results from a regression of HML_{t+1} onto future consumption growth; and cols. 4 and 5 use SMB_{t+1} in place of HML_{t+1} . Standard errors (in parentheses) are calculated under the null that the coefficients are zero, allowing for heteroskedasticity of returns. Col. 6 shows the test statistic and p -value (in brackets) for the test that the coefficients in the second and third regressions are zero, under the assumption that HML and SMB are uncorrelated, an assumption not rejected by the series. Col. 7 reports the variance of 100 times consumption growth.

To begin, we examine the predictive power of these size and value excess returns for current and future consumption growth. Column 1 of table 4 shows the fit of a regression of contemporaneous and future consumption growth, $\ln(C_{t+1+s}/C_t)$, onto HML_{t+1} and SMB_{t+1} for different S . The two excess returns jointly explain 1.78 percent of the variation in contemporaneous consumption growth ($S = 0$). As we increase S , the variance of the dependent variable increases (col. 7 of table 4). Despite this, the fit of the regression nearly doubles over the first four quarters. That is, over the first year, the additional future

consumption contains more “signal”—movement of consumption correlated with past returns—than “noise”—innovations to consumption not predicted by either excess return. The share of consumption growth explained by the size and value excess returns is above the $S = 0$ level out to three years.

To formally test the null hypothesis that there is no correlation between these portfolios and future consumption growth, we run two separate “reverse” regressions: regressing first HML_{t+1} and then SMB_{t+1} onto only *future* consumption growth, $\ln(C_{t+1+S}/C_{t+1})$. This specification makes inference simpler and sharper: under the null that there is no correlation between the return and future consumption growth, the residual of the regression inherits the time-series properties of the returns and so has no significant serial correlation. Columns 2–5 of table 4 show that the coefficients on future consumption growth display different patterns with S for HML and SMB, but both coefficients are typically positive. For the regression using HML_{t+1} , future consumption growth is statistically significant only for the horizon covering the three years after the return ($t + 1$ to $t + 1 + 11$). For SMB, future consumption growth is not statistically significant at any horizon, although it is close at horizons around one year.¹⁵

In terms of joint significance (col. 6), the significance level lies between 15 and 30 percent for $S = 3, \dots, 9$, then drops to the 10 percent level and then the 5 percent level at $S = 11$, and rises significantly thereafter. While these results may seem surprising to those who think of consumption as close to a martingale, several papers have documented that past aggregate returns predict consumption growth, from Hall (1978) to the subsequent analyses of Fama (1981), Daniel and Marshall (1997), and the papers cited in the Introduction.¹⁶

While table 4 focuses on testing, figure 2 presents point estimates. Figures 2*a* and 2*b* display the partial correlations of HML_{t+1} and SMB_{t+1} , respectively, with $\Delta \ln C_{t+1+S}$ and two standard error bands. Relative to the contemporaneous correlations, there is significant correlation between these excess returns and future consumption growth. The contemporaneous correlation between HML and $\Delta \ln C$ is very close

¹⁵ At its peak, future consumption growth explains 1.47 percent of the variation in HML and 2.23 percent of the variation in SMB, relative to 0.15 percent and 1.11 percent, respectively, at $S = 0$. These shares of variance explained by consumption growth are relatively large. Lettau and Ludvigson (2001*a*) demonstrate that a variable based on the ratio of consumption to wealth, denoted *cay* (see Sec. VII), predicts aggregate excess returns well. For comparison, *cay* explains only 0.5 percent of the variance of HML and 0.3 percent of SMB.

¹⁶ Consumption growth also does not appear to be a univariate martingale difference sequence. If consumption growth were a martingale, the variance in col. 7 of table 4 would increase linearly in S , which it does not. Over the first year, the increases in variance are 0.71, 0.87, 1.00, and 0.90, all greater than the initial variance of 0.64. But starting at two years, the increases are smaller: 0.49, 0.51, 0.39, 0.55, 0.64, and 0.60 for $S = 7, \dots, 13$.




FIG. 2.—Excess returns and future consumption growth: *a*, correlation between HML_{t+1} and $\ln(C_{t+1+S}/C_{t+S})$; *b*, correlation between SMB_{t+1} and $\ln(C_{t+1+S}/C_{t+S})$; *c*, covariance between HML_{t+1} and $\ln(C_{t+1+S}/C_t)$; *d*, covariance between SMB_{t+1} and $\ln(C_{t+1+S}/C_t)$. Two standard error bands are calculated using Newey-West standard errors with lags equal to $S+1$, since we are estimating under the assumption that consumption growth is predictable.

to zero, economically and statistically. The correlation between HML and future consumption growth rates, however, is typically positive and larger, until $S = 12$, and is statistically significant at $S = 8$. The correlation between SMB and contemporaneous consumption growth is significantly larger than for HML, and the correlation remains positive until $S = 8$.

Figures 2*c* and *d* display the cumulative covariances between consumption growth and each excess return, which are the covariances on the right-hand side of equation (8). The covariance of long-horizon consumption growth and HML (fig. 2*c*) is roughly zero until $S = 6$, and at $S = 11$ the covariance becomes borderline significant. For SMB (fig.

2d), the covariance rises with the horizon over the first two years and then declines.

We can use equation (8) and the covariances of figure 2 to quantify the impact of including future consumption growth in consumption risk. The mean of HML is 0.81 percent per quarter and the mean of SMB is 0.39 percent per quarter. Thus the contemporaneous consumption covariances imply a coefficient of relative risk aversion that is 69.5 for SMB and one that is negative for HML (although very close to a large positive number). When future consumption growth is added, by $S = 6$, both measures are positive and risk aversion based on SMB is 18.6, nearly one-fourth that estimated by contemporaneous consumption risk. At a horizon of $S = 11$, the implied levels of risk aversion are 45.0 and 25.6, respectively.

The fact that future consumption growth is predictable by HML and SMB makes it unlikely that our main findings presented in the previous section are spurious. That is, table 4 and figure 2 demonstrate that the differences in ultimate consumption risk across size and value portfolios are statistically significant. Table 1 does not merely relate average returns to covariances that have a pattern of noise that happens to help the model fit the data. As complementary evidence, Parker and Julliard (2003) show in their figure 2 that consumption betas at $S = 11$ have a greater spread and are more closely related to average returns than at $S = 0$.

We checked this concern further with a Monte Carlo experiment, drawing random sample paths of consumption growth and using them to price the Fama-French portfolios. Keeping the returns data, we generate 1,820 artificial data sets by randomly selecting consumption growth rates for 1 to T from the observed sample of consumption growth rates. We then estimate our model for $S = 0, \dots, 15$ on each generated data set using GMM with the prespecified weighting matrix.

First, how likely would we be to find a fit like that in table 1 if consumption growth were truly independent over time? On the basis of our simulation, the probability that $\gamma_s > 0$ and the fit is 44 percent or larger (as in table 1) for any consumption model ($S \in [0, 15]$) is one in five; for a fit of 60 percent or larger (table 3), the probability is 4.4 percent. But we also find $\gamma_s > 0$ for all S and a hump-shaped pattern of fit. The probability that we find $\gamma_s > 0$ for all S and $R^2 \geq 44$ percent (60 percent) for some S is 6.5 percent (1.8 percent).

Second, how likely would we be to find an increases in fit like that in table 1 if consumption growth were truly independent over time? To answer this question, we discard samples in which $\hat{\gamma} \leq 0$ for $S = 0$, leaving 977 samples in which the sign of the contemporaneous correlation between returns and consumption growth matches that in the actual data. Only 4.91 percent of the simulations estimate $\hat{\gamma}_s > 0$ for all S and




FIG. 3.—Distribution of $R_S^2 - R_0^2$ when consumption is a martingale. Sample of simulations in which risk aversion at $S = 0$ is positive and the R^2 lies between 0.026 and 0.244 at $S = 0$. At any horizon $S > 0$, R^2 is set to zero if risk aversion is estimated to be negative.

an increase in fit of 41 percent or more from $S = 0$ to any $S \in [1, 15]$. To check that the presence of some large R^2 's at $S = 0$ is not driving these results, we limit the R^2 's at $S = 0$ to a range similar to that of tables 1 and 3 by dropping simulations in the top and bottom 25 percent of the distribution of fit at $S = 0$, so that all simulations have $R^2 \in [0.026, 0.244]$ at $S = 0$. Figure 3 shows that in these simulations, there is no tendency for the consumption model to fit expected returns better as the horizon increases. In these simulations, the probability that $\hat{\gamma}_S > 0$ for all S and R_S^2 increases by at least 41 percent from $S = 0$ to any $S \in [1, 15]$ is 7.53 percent (36 out of 478 simulations), or 3.96 percent if we do not restrict ourselves to cases with $\gamma_0 > 0$. It is also the case that the simulated results show little correlation in R^2 across horizon, which is not penalized in these probabilities. We conclude that while of course our results could be due to chance, this is not likely.¹⁷

The serial correlation of consumption growth has a strong business cycle pattern, and it is this that prices expected returns. Figure 4 displays the estimated stochastic discount factors for $S = 0$, for $S = 11$, and for $S = 11$ with the risk-free rate constant. The figure also shows, as shaded

¹⁷ Parker and Julliard (2003) reports another check on the results. Kan and Zhang (1999), with a simulation exercise, shows that “useless” factors can appear statistically significant, at least when the Fama and MacBeth (1973) methodology is applied to a misspecified model. If a factor is spurious, its significance tends to be reduced by firm characteristics (Jagannathan and Wang 1996, 1998). When Fama-MacBeth estimation and a linearized model are used, ultimate consumption risk remains significantly related to returns, and more significantly than contemporaneous consumption risk, when we include as factors two different asset characteristics—size and book-to-market value—and these factors are insignificant in three out of four specifications.




FIG. 4.—Time series of stochastic discount factors. The dashed line displays m_{t+1}^0 , the solid line displays m_{t+1}^{11} , and the long-dashed line (lying almost completely under the solid line) displays m_{t+1}^{11} with $R^l = 1$. Risk aversion $\hat{\gamma}$'s are set to the values reported in table 1, first set of results. Shaded regions are NBER recessions.

regions, the NBER recession dates. While consumption growth at $S = 0$ has little visible business cycle pattern, at $S = 11$, the series is clearly related to the business cycle. Consumption falls around recessions, so that the stochastic discount factor is highest right before and at the start of recessions. Ultimate consumption risk prices expected returns on the Fama-French portfolios because value stocks, and to some extent small stocks, pay off relatively poorly before and at the start of recessions, when consumption enters a period of decline.

Figure 4 also makes clear the small role played by the risk-free real interest rate in our stochastic discount factor. That is, our emphasis in this section on consumption growth alone is not misleading. This point also comes through in the next section and, in more detail, in Parker and Julliard (2003).

To conclude this section, can one choose a “best” horizon of consumption adjustment at which to measure ultimate consumption risk? On the one hand, if consumption is slow to adjust for the reasons outlined above, then larger S 's are preferred since they allow a longer horizon for consumption adjustment. The fact that the economic importance of consumption risk varies with horizon is evidence against the basic CCAPM used in Section II and evidence in favor of a modified model in which consumption exhibits slow adjustment for some reason. This suggests that one would like to choose S as large as possible. On the other hand, as S increases, the noise in m_{t+1}^S increases. Noise in a factor is irrelevant asymptotically, but it decreases the accuracy of estimation in any finite sample. For larger S , the noise in m_{t+1}^S is greater because we do not observe $E_{t+1}[m_{t+1}^S]$, and all innovations to consumption between $t + 1$ and $t + 1 + S$ are included as noise in our stochastic discount factor.¹⁸ Closely related, for larger S , our sample is effectively shorter because there is more correlation over time in GMM residuals, leading to less precision in estimated covariances.¹⁹ Thus we do not want to choose S too large.

These arguments suggest that the best horizon has a large signal and little noise in the measured covariance between the discount factor and returns. As discussed above, column 6 of table 4 speaks to this issue. The statistical significance of the covariances of HML and SMB with long-horizon consumption growth peaks at $S = 11$ and declines sub-

¹⁸ Any construction of $E_{t+1}[m_{t+1}^S]$ would still contain noise and, potentially more importantly, would omit a signal—true slow adjustment of consumption—from our factor.

¹⁹ 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 implies a shorter time series of available returns to price.

stantially thereafter.²⁰ Parker and Julliard (2003) analyze all 25 portfolios and reach a similar conclusion. They compute, for each S , the average statistical uncertainty in consumption betas estimated from a linearized ultimate consumption risk model ($\hat{\beta}_{i,S}$) relative to the total cross-sectional variance of consumption betas:

$$\frac{(1/25) \sum_{i=1}^{25} \widehat{\text{Var}}(\hat{\beta}_{i,S})}{(1/25) \sum_{i=1}^{25} [\hat{\beta}_{i,S} - (1/25) \sum_{i=1}^{25} \hat{\beta}_{i,S}]^2}.$$

The numerator is a measure of the average noise in $\hat{\beta}_{i,S}$, and the denominator is a measure of the total observed signal plus noise across $\hat{\beta}_{i,S}$. Statistical uncertainty in the estimated $\beta_{i,S}$ remains roughly constant relative to the total variance from $S = 0$ to $S = 11$ and is substantially larger for larger S . On these grounds, one might choose to measure ultimate consumption risk in this sample at $S = 11$.

But tables 1, 3, and 5 (below) show that for S close to three years, the conclusions one reaches are quite similar. 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, but typically $S = 11$ maximizes the model's fit.

VII. Comparison with Other Linear Factor Models

In this section, we compare the performance of a linear version of our ultimate consumption risk model to the linear asset pricing models of Fama and French and Lettau and Ludvigson.

Fama and French (1992, 1993) show that a three-factor model explains a large fraction of the cross-sectional variation in expected returns in the Fama-French portfolios. The factors are the excess return on the market (denoted R^m) and the two excess returns capturing the size and value premia already discussed, HML and SMB. Lettau and Ludvigson (2001*a*) argues that the budget constraint of the representative household implies that consumption, income, and asset 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 (2001*b*) shows that this variable, denoted by $ca y_p$ consumption growth ($\Delta \ln C_{t+1}$), and their interaction provide a

²⁰ This is also true in a regression of $\ln(C_{t+1+S}/C_t)$ onto both returns, correcting for serial correlation. Figure 2, however, shows that this pattern differs by size and value, and on the basis of the correlations between consumption and SMB alone, one would choose a lower "best" S .

three-factor model that does as well in explaining the cross section of expected returns as the Fama-French three-factor model.

To make ultimate consumption risk into a linear model comparable to these models, we follow Lettau and Ludvigson and apply a first-order log-linear approximation to the utility function to yield

$$m_t^S = R_{t,t+S}^f - \gamma_S R_{t,t+S}^f \ln \left(\frac{C_{t+S}}{C_{t-1}} \right).$$

This is then a two-factor model with factors $R_{t,t+S}^f$ and $R_{t,t+S}^f \times \ln(C_{t+S}/C_{t-1})$. We do not impose the model's restriction on the coefficient on $R_{t,t+S}^f$. However, to emphasize that the covariance of returns and consumption growth is pricing expected returns and that the risk-free interest rate is doing very little, we also report results that maintain the assumption that the risk-free real interest rate is constant, making our model a one-factor model in which the stochastic discount factor is simply long-horizon consumption growth.

Each model says that the expected return on any portfolio is the weighted sum of the covariance of the return and each factor. Denote the vector of factors by f_{t+1} , so

$$f_{t+1} = (cay_t, \Delta \ln C_{t+1}, cay_t \Delta \ln C_{t+1})'$$

in the Lettau-Ludvigson model,

$$f_{t+1} = (R_{t+1}^m, \text{SMB}_{t+1}, \text{HML}_{t+1})'$$

in the Fama-French model, and

$$f_{t+1} = \left(R_{t+1,t+1+S}^f \ln \frac{C_{t+1+S}}{C_t}, R_{t+1,t+1+S}^f, 0 \right)'$$

in the contemporaneous and ultimate consumption risk models. Let $\mathbf{b} = (b_1, b_2, b_3)'$ be the vector of coefficients on the factors. Following Yogo (2003), we estimate the Fama-French and Lettau-Ludvigson models by GMM, using the 28×1 empirical moment function

$$g(\mathbf{R}_t^e, f_{t+1}; \alpha, \mu, \mathbf{b}) = \begin{bmatrix} \mathbf{R}_t^e - \alpha \mathbf{1}_{25} + \mathbf{R}_t^e (f_t - \mu)' \mathbf{b} \\ f_t - \mu \end{bmatrix},$$

where μ now denotes a 3×1 parameter vector. Under the null that the model prices expected returns, the theoretical moment restriction $E[g(\mathbf{R}_t^e, f_{t+1}; \alpha, \mu, \mathbf{b})] = 0$ holds for the true $(\alpha, \mu', \mathbf{b}') \in \mathbb{R}^7$. As in our basic estimation, the difference between a fitted moment and zero is a measure of the mispricing of an expected return, and we include an intercept that allows all excess returns to be mispriced by a common amount. Finally, we estimate the consumption risk model with time-varying R^f by imposing $\mu_3 = b_3 = 0$ and omitting the last moment; we

estimate the model with a constant R^f by imposing $\mu_2 = \mu_3 = b_2 = b_3 = 0$ and omitting the last two moments.²¹

We present results for the subsample analyzed by Lettau and Ludvigson (2001*b*) and the focus of Fama and French (1993).²² The first row of panel A of table 5 reports the fit, estimated intercept and coefficients, and Hansen-Jagannathan distance and p -value for the Fama-French three-factor model.²³ The second row of results reports the same set of statistics for the Lettau-Ludvigson three-factor model. The remaining rows report results for the contemporaneous CCAPM and the ultimate consumption risk model for $S = 9, \dots, 14$ quarters. For these models, we also present the implied coefficient of relative risk aversion: $\hat{\gamma}_s = \hat{b}_1 \hat{\mu}_{2,s} (1 + \hat{b}_1 \hat{\mu}_{1,s})^{-1}$ for the model with time-varying R^f and $\hat{\gamma}_s = \hat{b}_1 (1 + \hat{b}_1 \hat{\mu}_s)^{-1}$ for the model without.

The two main points of table 5 are (1) that the ultimate consumption risk model with R^f constant, a single-factor model, fits expected returns nearly as well as the three-factor models of Lettau and Ludvigson and Fama and French; and (2) that the ultimate consumption risk two-factor model actually fits expected returns slightly better than both models.²⁴ The explanatory power of the Fama-French, Lettau-Ludvigson, and ultimate risk to consumption models are all economically significant, fitting 67 percent, 64 percent, and 55–70 percent of the variation in expected returns, respectively.²⁵ The contemporaneous CCAPM performs poorly.

Long-horizon consumption growth also implies lower levels of the estimated intercept; the Lettau-Ludvigson and Fama-French models perform less well on this dimension. Finally, the ultimate consumption risk model fits with quite low levels of estimated risk aversion. Efficient GMM estimates are all well below 10.

Figure 5 graphs the pricing errors for each portfolio, for the four

²¹ As a prespecified weighting matrix, we use an identity matrix, resetting the diagonal entries for the moments $E[f_t - \mu] = 0$ to very large numbers so that the point estimates are identical to those from the Fama and MacBeth (1973) procedure.

²² Parker and Julliard (2003) presents similar results for the longest sample in which data for all three models are available.

²³ Covariance matrices are calculated using the Newey-West procedure with one lag for the three-factor models and $S + 1$ lags for the ultimate consumption risk model. Given the potential unreliability of nonparametric covariance matrix estimation in a small sample, we also experimented with alternative approaches: Newey-West standard errors computed with 25 lags, Hansen-Hodrick standard errors with $S + 1$ lags and no downweighting, and VARHAC covariance matrix estimation. All these procedures delivered standard errors numerically roughly within 10 percent of the ones reported in table 5.

²⁴ The coefficient on the risk-free real interest rate (b_2) should be one over the mean of the stochastic discount factor, $1/(\bar{\mu}_2 - \gamma \bar{\mu}_1)$. This restriction is not rejected for any model with $S > 9$. It is rejected for $S = 2, \dots, 5$ with the prespecified weighting matrix and for $S = 2, \dots, 9$ with efficient GMM.

²⁵ In some of these cases, the Hansen-Jagannathan distance measure does not approximate the average pricing error.

TABLE 5
COMPARISON OF AFFINE FACTOR MODELS OF EXPECTED RETURNS

MODEL	GMM WITH PRESPECIFIED WEIGHTING MATRIX							EFFICIENT GMM					
	R^2	α	b_1	b_2	b_3	γ	Dist	α	b_1	b_2	b_3	γ	JTest
A. Lettau-Ludvigson and Fama-French Three-Factor Models and Contemporaneous CCAPM													
FF	.67	.013 (.014)	1.46 (3.47)	1.14 (2.42)	4.18 (2.74)		.50 [.0000]	.021 (.010)	.35 (2.46)	1.73 (2.42)	4.19 (2.09)		42.1 [.0041]
LL	.64	.019 (.010)	-5.3 (33.7)	47.5 (21.8)	5,642 (3,196)		.55 [.0000]	.006 (.006)	-10.6 (22.9)	81.3 (21.8)	-733.9 (1,831.0)		40.9 [.0057]
S=0*	.00	.000 (.000)	.0 (41.4)			-3.0 (42.3)	2.93 [.0000]	.000 (.000)	.81 (23.96)			.80 (23.81)	49.9 [.0010]
S=0	.18	.016 (.007)	59.8 (55.7)			48.0 (35.8)	1.30 [.0000]	.020 (.006)	89.74 (22.51)			65.81 (11.98)	42.4 [.0082]
B. Ultimate Consumption Risk Model with Time-Varying R^f													
S=9	.63	.005 (.008)	32.5 (23.1)	3.2 (5.1)		13.7 (3.7)	.53 [.0000]	.008 (.005)	13.59 (7.57)	11.30 (5.10)		8.58 (2.75)	35.8 [.0320]
S=10	.65	.004 (.009)	31.7 (22.0)	2.9 (5.0)		12.8 (3.3)	.52 [.0001]	.008 (.005)	11.94 (7.02)	9.95 (4.98)		7.62 (2.64)	36.5 [.0271]
S=11	.70	.004 (.011)	33.2 (19.6)	3.0 (4.7)		12.4 (2.5)	.48 [.0002]	.009 (.005)	14.69 (6.71)	8.10 (4.69)		8.40 (1.90)	37.0 [.0237]
S=12	.68	-.001 (.013)	36.8 (25.3)	3.5 (4.5)		12.2 (2.6)	.50 [.0004]	.010 (.005)	11.62 (6.65)	7.74 (4.50)		7.05 (2.21)	37.2 [.0223]
S=13	.65	-.002 (.015)	43.0 (32.1)	2.0 (4.4)		12.2 (2.5)	.51 [.0014]	.011 (.004)	10.34 (7.26)	6.59 (4.37)		6.36 (2.54)	38.7 [.0154]
S=14	.62	-.001 (.015)	43.2 (32.1)	1.9 (4.2)		11.6 (2.3)	.54 [.0013]	.013 (.005)	8.72 (7.37)	6.07 (4.18)		5.57 (2.85)	38.7 [.0153]

C. Ultimate Consumption Risk Model with Constant R^f									
$S=9$	.57	.007	32.7		14.2	1.4	.019	5.14	4.27 50.86
		(.009)	(23.5)		(4.0)	[.0000]	(.006)	(7.82)	(5.34) [.0007]
$S=10$	.59	.006	32.2		13.4	1.4	.019	4.45	3.73 50.80
		(.009)	(22.6)		(3.5)	[.0000]	(.006)	(7.63)	(5.31) [.0007]
$S=11$	.66	.005	35.7		13.3	1.4	.017	8.69	6.20 51.99
		(.011)	(20.7)		(2.5)	[.0001]	(.006)	(7.06)	(3.41) [.0005]
$S=12$	.61	.001	40.0		13.2	1.5	.017	8.73	6.08 52.69
		(.013)	(26.1)		(2.6)	[.0002]	(.006)	(7.39)	(3.40) [.0004]
$S=13$	.59	.000	47.3		13.3	1.4	.017	7.02	5.09 51.65
		(.015)	(35.8)		(2.6)	[.0003]	(.006)	(7.96)	(4.05) [.0006]
$S=14$	.55	.000	48.0		12.7	1.4	.017	7.94	5.45 51.97
		(.015)	(36.7)		(2.4)	[.0004]	(.006)	(7.87)	(3.55) [.0005]

NOTE.—GMM estimation of affine factor models using returns data from 1963:3–1999:4. Standard errors are reported in parentheses and p -values in brackets. GMM with a prespecified weighting matrix replicates the Fama-MacBeth point estimates by using an identity matrix for the moments corresponding to expected returns and large weights on the diagonal for the remaining moments, as discussed in Sec. IV. Efficient GMM iterates to convergence. The J -statistics for the Fama-French and Lettau-Ludvigson models are distributed according to a $\chi^2(21)$ under null, whereas the distributions for the consumption models are $\chi^2(22)$ and $\chi^2(23)$. Covariance matrices are calculated using the Newey-West procedure with one lag for the three-factor models and $S+1$ for the ultimate consumption risk model.

* Uses the beginning-of-period timing convention. Denotes using the timing convention for consumption data that consumption takes place at the start of the period.




FIG. 5.—Comparison of affine factor models of expected returns: *a*, contemporaneous CCAPM; *b*, Fama-French three-factor model; *c*, Lettau-Ludvigson *cay* model; *d*, ultimate consumption risk, $S = 11$. All returns are quarterly rates. Each portfolio is denoted by the rank of its market equity and then the rank of its ratio of book value to market value. Fitted values are based on the model estimates from panels A and B of table 5, which use returns data from 1963:3–1999:4.

main models. All models besides the contemporaneous CCAPM do quite well at fitting expected returns.

VIII. Conclusion

This paper analyzes the Fama-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 equilibrium risk of each marginal investment to the utility of investors. We show that while the covariance of each portfolio and contemporaneous consumption growth does not explain the pattern of average returns across portfolios, the ultimate risk to consumption explains a large fraction of the variation in average returns. The fit of our model recast as a linear one-factor model rivals that of the three-factor model of Fama and French and that of the three-factor model of Lettau and Ludvigson. These conclusions are robust to several variations in assumptions.

In sum, this paper confirms one of the central insights of the CCAPM: that consumption risk is an important determinant of average returns across stocks. This finding raises several questions. First, does the ultimate risk to consumption explain differences in expected bond returns or differences in expected returns over time? Parker (2003) shows that contemporaneous consumption risk is negatively related to time variation in expected returns and ultimate consumption risk is positively related. Second, estimates of the risk aversion of the representative household are still larger than we find plausible. Parker (2001) uses data on the consumption of households that actually hold stock and shows that the ultimate consumption risk of stockholders and the average premium on equity imply levels of risk aversion less than 10. Thus it may be the case that the ultimate consumption risk of different portfolios for stockholders provides an even better accounting of the cross section of expected returns.

Appendix

The Hansen-Jagannathan Distance

This Appendix extends the Hansen-Jagannathan distance test to the case of moment conditions that are nonlinear in the parameters, by extending theorem 3 of Jagannathan and Wang (1996). Let $\theta = (\mu_s, \gamma_s, \alpha_s)$ and $w_t = (R_t^e, C_{t+s}, C_{t-1})'$. The Hansen-Jagannathan distance and its sample analogue are given by

$$\begin{aligned} \text{dist}(w_p, \theta) &= \sqrt{E[g(w_p; \theta)]' \hat{W} E[g(w_p; \theta)]}, \\ \text{dist}_T(w_p, \theta) &= \sqrt{\min_{\theta} g_T(w_p; \theta)' \hat{W}_{g_T} g_T(w_p; \theta)}, \end{aligned} \quad (\text{A1})$$

where

$$g_T(\mathbf{w}_p, \boldsymbol{\theta}) = \frac{1}{T} \sum_{t=0}^T g(\mathbf{w}_p, \boldsymbol{\theta}).$$

If the weighting matrix $\hat{\mathbf{W}}$ is optimal in the sense of Hansen (1982), then $T[\text{dist}_T(\mathbf{w}_p, \boldsymbol{\theta})]^2$ is asymptotically distributed as a $\chi^2(M-P)$, where M is the number of moment restrictions and P the number of parameters. For any positive definite matrix $\mathbf{W} \xrightarrow{P} \hat{\mathbf{W}}$ that is not the efficient weighting matrix, the distribution of the test statistics is nonstandard.

PROPOSITION. Let $\{\mathbf{w}_t\}$ be ergodic stationary and let $\hat{\boldsymbol{\theta}}$ be the GMM estimator defined as the minimizer of equation (A1), where $\mathbf{W}$ converges in probability to some symmetric positive definite matrix $\mathbf{W}$. Suppose that the model is correctly specified in that $E[g(\mathbf{w}_p, \boldsymbol{\theta}_0)] = \mathbf{0}_M$ holds for some $\boldsymbol{\theta}_0 \in \Theta$. Suppose that (i) the parameter space Θ is a compact subset of $\mathbb{R}^p$, (ii) $g(\mathbf{w}_p, \boldsymbol{\theta})$ is continuous in $\boldsymbol{\theta}$ for all $\mathbf{w}_p$, (iii) $g(\mathbf{w}_p, \boldsymbol{\theta})$ is measurable in $\mathbf{w}_t$ for all $\boldsymbol{\theta} \in \Theta$, (iv) $E[g(\mathbf{w}_p, \boldsymbol{\theta})] \neq \mathbf{0}_M$ for all $\boldsymbol{\theta} \neq \boldsymbol{\theta}_0$ in Θ , and (v) $E[\sup_{\boldsymbol{\theta} \in \Theta} \|g(\mathbf{w}_p, \boldsymbol{\theta})\|] < \infty$. Suppose further that

1. $\boldsymbol{\theta}_0$ is in the interior of Θ ;
2. $g(\mathbf{w}_p, \boldsymbol{\theta})$ is continuously differentiable in $\boldsymbol{\theta}$ for any $\mathbf{w}_p$;
- 3.

$$\sqrt{T}g_T(\mathbf{w}_p, \boldsymbol{\theta}_0) \xrightarrow{d} N(\mathbf{0}_M, \boldsymbol{\Omega}_{(M \times M)}),$$

4. where $\boldsymbol{\Omega}$ is positive definite;
5. (local dominance condition on $\partial g(\mathbf{w}_p, \boldsymbol{\theta})/\partial \boldsymbol{\theta}'$)

$$E \left[\sup_{\boldsymbol{\theta} \in \mathcal{N}} \left\| \frac{\partial g(\mathbf{w}_p, \boldsymbol{\theta})}{\partial \boldsymbol{\theta}'} \right\| \right] < \infty$$

6. for some neighborhood $\mathcal{N}$ of $\boldsymbol{\theta}_0$; and
- 7.

$$\mathbf{G}_{(M \times P)} = E \left[\frac{\partial g(\mathbf{w}_p, \boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta}'} \right]$$

is of full column rank.

Let

$$\mathbf{A} = \boldsymbol{\Omega}^{1/2} \mathbf{W}^{1/2} [\mathbf{I}_M - (\mathbf{W}^{1/2})' \mathbf{G} (\mathbf{G}' \mathbf{W} \mathbf{G})^{-1} \mathbf{G}' \mathbf{W}^{1/2}] (\mathbf{W}^{1/2})' (\boldsymbol{\Omega}^{1/2})', \quad (\text{A2})$$

where $\boldsymbol{\Omega}^{1/2}$ and $\mathbf{W}^{1/2}$ are the upper-triangular matrices from the Choleski decomposition of $\boldsymbol{\Omega}$ and $\mathbf{W}$, and $\mathbf{I}_M$ is an M -dimensional identity matrix. Then, $\mathbf{A}$ has exactly $M-P$ nonzero eigenvalues, which are positive and denoted by $\lambda_1, \dots, \lambda_{M-P}$ and the asymptotic sampling distribution of the Hansen-Jagannathan distance is

$$T[\text{dist}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})]^2 \xrightarrow{d} \sum_{j=1}^{M-P} \lambda_j v_j^2 \quad \text{as } T \rightarrow \infty,$$

where $v_1, \dots, v_{M-P}$ are independent $\chi^2(1)$ random variables.

Proof. The first-order condition of the minimization problem is

$$\mathbf{G}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})' \hat{\mathbf{W}} g_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}}) = \mathbf{0}_P, \quad (\text{A3})$$

where

$$\mathbf{G}_T(\mathbf{w}_p; \hat{\boldsymbol{\theta}}) = \frac{1}{T} \sum_{t=0}^T \frac{\partial g(\mathbf{w}_p, \hat{\boldsymbol{\theta}})}{\partial \boldsymbol{\theta}'}$$

Since $g(\mathbf{w}_p, \boldsymbol{\theta})$ is continuously differentiable, we can apply the mean value theorem

$$g_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}}) = g_T(\mathbf{w}_p, \bar{\boldsymbol{\theta}}) + \mathbf{G}_T(\mathbf{w}_p; \bar{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}} - \bar{\boldsymbol{\theta}}), \quad (\text{A4})$$

where $\bar{\boldsymbol{\theta}}$ is a mean value lying between $\hat{\boldsymbol{\theta}}$ and $\boldsymbol{\theta}_0$, and substitute into equation (A3)

$$\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 = -[\mathbf{G}_T(\mathbf{w}_p, \bar{\boldsymbol{\theta}})' \hat{\mathbf{W}} \mathbf{G}_T(\mathbf{w}_p; \bar{\boldsymbol{\theta}})]^{-1} \mathbf{G}_T(\mathbf{w}_p; \hat{\boldsymbol{\theta}}) \hat{\mathbf{W}} g_T(\mathbf{w}_p, \boldsymbol{\theta}_0).$$

From equation (A4), the sample analogue of the moment function is

$$g_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}}) = [\mathbf{I}_M - \mathbf{G}_T(\mathbf{w}_p; \bar{\boldsymbol{\theta}})[\mathbf{G}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})' \hat{\mathbf{W}} \mathbf{G}_T(\mathbf{w}_p; \bar{\boldsymbol{\theta}})]^{-1} \mathbf{G}_T(\mathbf{w}_p; \hat{\boldsymbol{\theta}}) \hat{\mathbf{W}} g_T(\mathbf{w}_p, \boldsymbol{\theta}_0),$$

which, in the definition of the Hansen-Jagannathan distance, gives

$$T[\text{dist}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})]^2 \xrightarrow{d} \mathbf{Z}'[\mathbf{W} - \mathbf{W}\mathbf{G}(\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}\mathbf{G}'\mathbf{W}]\mathbf{Z} \quad \text{as } T \rightarrow \infty, \quad (\text{A5})$$

where $\mathbf{Z}$ is an M -dimensional vector of mean zero normal random variables, $\mathbf{Z} \sim N(\mathbf{0}_M, \boldsymbol{\Omega})$. Let $\mathbf{z} \sim N(\mathbf{0}_M, \mathbf{I}_M)$ so that $\mathbf{Z} = (\boldsymbol{\Omega}^{1/2})\mathbf{z}$. Substituting into equation (A5) yields

$$T[\text{dist}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})]^2 \xrightarrow{d} \mathbf{z}'\mathbf{A}\mathbf{z},$$

where $\mathbf{A}$ is defined in equation (A2). Since

$$\mathbf{I}_M - (\mathbf{W}^{1/2})'\mathbf{G}(\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}\mathbf{G}'\mathbf{W}^{1/2}$$

is symmetric and idempotent and has trace $M - P$, we know that it has rank $M - P$. This implies that the rank of $\mathbf{A}$ is also $M - P$. As a consequence, $\mathbf{A}$ has exactly $M - P$ positive eigenvalues, denoted by $\lambda_1, \dots, \lambda_{M-P}$. Then, there exist a diagonal matrix $\boldsymbol{\Lambda}$ and an orthogonal matrix $\mathbf{H}$ such that $\mathbf{A} = \mathbf{H}\boldsymbol{\Lambda}\mathbf{H}$, where $\boldsymbol{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_{M-P}, \mathbf{0}_P\}$ and $\mathbf{0}_P$ is a P -dimensional vector of zeros. Let $\mathbf{x} = \mathbf{H}\mathbf{z}$; then $\mathbf{x} \sim N(\mathbf{0}_M, \mathbf{I}_M)$, and we have

$$T[\text{dist}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})]^2 \xrightarrow{d} \mathbf{x}'\mathbf{A}\mathbf{x} = \sum_{j=1}^{M-P} \lambda_j x_j^2.$$

Setting $v_j = x_j^2$ completes the proof.

We estimate the matrix $\mathbf{A}$ by

$$\hat{\mathbf{A}} = \hat{\boldsymbol{\Omega}}^{1/2} \hat{\mathbf{W}}^{1/2} [\mathbf{I}_M - (\hat{\mathbf{W}}^{1/2})'\hat{\mathbf{G}}(\hat{\mathbf{G}}'\hat{\mathbf{W}}\hat{\mathbf{G}})^{-1}\hat{\mathbf{G}}'\hat{\mathbf{W}}^{1/2}] (\hat{\mathbf{W}}^{1/2})' (\hat{\boldsymbol{\Omega}}^{1/2})',$$

where $\hat{\mathbf{G}} = \mathbf{G}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})$ and $\hat{\boldsymbol{\Omega}}$ is a consistent estimate of $\boldsymbol{\Omega}$. We estimate $\lambda_1, \dots, \lambda_{M-P}$ by the positive eigenvalues of $\hat{\mathbf{A}}$. Although $T[\text{dist}_T(\mathbf{w}_p, \hat{\boldsymbol{\theta}})]^2$ has a nonstandard distribution function, we can still consistently compute the p -values of the computed statistic to test the null hypothesis that the stochastic discount factor is correctly specified.

Define $u \equiv \sum_{j=1}^{M-P} \lambda_j v_j$, where u has the unknown probability distribution function $\psi(u)$. Let $\{v_{ij}\}$, $i = 1, \dots, T^*$, $j = 1, \dots, M - P$, be $T^*(M - P)$ independent random draws from a $\chi^2(1)$ distribution. We can construct a set of independent samples, $\{u_i\}_{i=1}^{T^*}$, by letting $u_i = \sum_{j=1}^{M-P} \lambda_j v_{ij}$. By the law of large numbers we have

that, for any nonnegative α ,

$$\frac{1}{T^*} \sum_{i=1}^{T^*} \mathbf{1}_{\{u_i \leq \alpha\}} \xrightarrow{p} \int_0^\alpha d\psi(u) = \Pr(u_i \leq \alpha) \quad \text{as } T^* \rightarrow \infty,$$

where $\mathbf{1}_{\{\cdot\}}$ is the index function that takes value one if the condition is satisfied and zero otherwise.

References

- Abel, Andrew B. 1990. "Asset Prices under Habit Formation and Catching Up with the Joneses." *A.E.R. Papers and Proc.* 80 (May): 38–42.
- Ahn, Seung C., and Christopher Gadarowski. 1999. "Small Sample Properties of the Model Specification Test Based on the Hansen-Jagannathan Distance." Manuscript, Dept. Econ., Arizona State Univ.
- Ait-Sahalia, Yacine, Jonathan A. Parker, and Motohiro Yogo. 2004. "Luxury Goods and the Equity Premium." *J. Finance* 59 (December): 2959–3004.
- Alvarez, Fernando, Andrew Atkeson, and Patrick J. Kehoe. 2002. "Money, Interest Rates, and Exchange Rates with Endogenously Segmented Markets." *J.P.E.* 110 (February): 73–112.
- Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad. Forthcoming. "Consumption, Dividends, and the Cross-Section of Equity Returns." *J. Finance*.
- Bansal, Ravi, and Amir Yaron. 2004. "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles." *J. Finance* 59 (August): 1481–1509.
- Brainard, William C., William R. Nelson, and Matthew D. Shapiro. 1991. "The Consumption Beta Explains Expected Returns at Long Horizons." Manuscript, Dept. Econ., Yale Univ.
- Breeden, Douglas T. 1979. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *J. Financial Econ.* 7 (September): 265–96.
- Breeden, Douglas T., Michael R. Gibbons, and Robert H. Litzenberger. 1989. "Empirical Tests of the Consumption-Oriented CAPM." *J. Finance* 44 (June): 231–62.
- Caballero, Ricardo J. 1995. "Near-Rationality, Heterogeneity, and Aggregate Consumption." *J. Money, Credit and Banking* 27 (February): 29–48.
- Campbell, John Y. 1996. "Understanding Risk and Return." *J.P.E.* 104 (April): 298–345.
- . 1999. "Asset Prices, Consumption, and the Business Cycle." In *Handbook of Macroeconomics*, vol. 1C, edited by John B. Taylor and Michael Woodford. Amsterdam: Elsevier Sci.
- Cochrane, John H. 1996. "A Cross-Sectional Test of an Investment-Based Asset Pricing Model." *J.P.E.* 104 (June): 572–621.
- . 2001. *Asset Pricing*. Princeton, NJ: Princeton Univ. Press.
- Daniel, Kent D., and David Marshall. 1997. "Equity-Premium and Risk-Free-Rate Puzzles at Long Horizons." *Macroeconomic Dynamics* 1 (2): 452–84.
- D'Avolio, Gene. 2002. "The Market for Borrowing Stock." *J. Financial Econ.* 66 (November–December): 271–306.
- Dynan, Karen E., and Dean M. Maki. 2001. "Does Stock Market Wealth Matter for Consumption?" Finance and Economics Discussion Series, no. 2001-23, Fed. Reserve Bd., Washington, DC.
- Eichenbaum, Martin S., Lars Peter Hansen, and Kenneth J. Singleton. 1988. "A

- Times Series Analysis of Representative Agent Models of Consumption and Leisure Choice under Uncertainty." *Q.J.E.* 103 (February): 51–78.
- Fama, Eugene F. 1981. "Stock Returns, Real Activity, Inflation, and Money." *A.E.R.* 71 (September): 545–65.
- Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *J. Finance* 47 (June): 427–65.
- . 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *J. Financial Econ.* 33 (February): 3–56.
- Fama, Eugene F., and James D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *J.P.E.* 81 (May/June): 607–36.
- Ferson, Wayne E., and Stephen R. Foerster. 1994. "Finite Sample Properties of the Generalized Method of Moments in Tests of Conditional Asset Pricing Models." *J. Financial Econ.* 36 (August): 29–55.
- Flavin, Marjorie. 2001. "Owner-Occupied Housing and the Presence of Adjustment Costs: Implications for Asset Pricing and Nondurable Consumption." Manuscript, Dept. Econ., Univ. California, San Diego.
- Gabaix, Xavier, and David Laibson. 2002. "The 6D Bias and the Equity-Premium Puzzle." In *NBER Macroeconomics Annual 2001*, vol. 16, edited by Ben S. Bernanke and Ken Rogoff. Cambridge, MA: MIT Press.
- Grossman, Sanford J., and Robert J. Shiller. 1981. "The Determinants of the Variability of Stock Market Prices." *A.E.R. Papers and Proc.* 71 (May): 222–27.
- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *J.P.E.* 86 (December): 971–87.
- Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (July): 1029–54.
- Hansen, Lars Peter, John Heaton, and Amir Yaron. 1996. "Finite-Sample Properties of Some Alternative GMM Estimators." *J. Bus. and Econ. Statis.* 14 (July): 262–80.
- Hansen, Lars Peter, and Ravi Jagannathan. 1997. "Assessing Specification Errors in Stochastic Discount Factor Models." *J. Finance* 52 (June): 557–90.
- Hansen, Lars Peter, and Kenneth J. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 50 (September): 1269–86.
- Hayashi, Fumio. 2000. *Econometrics*. Princeton, NJ: Princeton Univ. Press.
- Jagannathan, Ravi, and Zhenyu Wang. 1996. "The Conditional CAPM and the Cross-Section of Expected Returns." *J. Finance* 51 (March): 3–53.
- . 1998. "An Asymptotic Theory for Estimating Beta-Pricing Models Using Cross-Sectional Regression." *J. Finance* 53 (August): 1285–1309.
- Kan, Raymond M., and Chu Zhang. 1999. "Two-Pass Tests of Asset Pricing Models with Useless Factors." *J. Finance* 54 (February): 203–35.
- Kandel, Shmuel, and Robert F. Stambaugh. 1990. "Expectations and Volatility of Consumption and Asset Returns." *Rev. Financial Studies* 3 (2): 207–32.
- Lamont, Owen A., and Richard H. Thaler. 2003. "Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs." *J.P.E.* 111 (April): 227–68.
- Lettau, Martin, and Sydney Ludvigson. 2001a. "Consumption, Aggregate Wealth, and Expected Stock Returns." *J. Finance* 56 (June): 815–49.
- . 2001b. "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying." *J.P.E.* 109 (December): 1238–87.
- Ludvigson, Sydney, and Charles Steindel. 1999. "How Important Is the Stock Market Effect on Consumption?" *Fed. Reserve Bank New York Econ. Policy Rev.* 5 (July): 29–51.
- Lustig, Hanno, and Stijn Van Nieuwerburgh. 2003. "Housing Collateral, Con-

- sumption Insurance and Risk Premia: An Empirical Perspective." Working Paper no. 9959 (September), NBER, Cambridge, MA.
- Lynch, Anthony W. 1996. "Decision Frequency and Synchronization across Agents: Implications for Aggregate Consumption and Equity Return." *J. Finance* 51 (September): 1479–97.
- Mankiw, N. Gregory, and Matthew D. Shapiro. 1986. "Risk and Return: Consumption Beta versus Market Beta." *Rev. Econ. and Statis.* 68 (August): 452–59.
- Marshall, David A., and Nayan G. Parekh. 1999. "Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?" *J. Finance* 54 (April): 623–54.
- Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *J. Monetary Econ.* 15 (March): 145–61.
- Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (May): 703–8.
- Parker, Jonathan A. 2000. "Spendthrift in America? On Two Decades of Decline in the U.S. Savings Rate." In *NBER Macroeconomics Annual 1999*, vol. 14, edited by Ben S. Bernanke and Julio J. Rotemberg. Cambridge, MA: MIT Press.
- . 2001. "The Consumption Risk of the Stock Market." *Brookings Papers Econ. Activity*, no. 2, 279–333.
- . 2003. "Consumption Risk and Expected Stock Returns." *A.E.R. Papers and Proc.* 93 (May): 376–82.
- Parker, Jonathan A., and Christian Julliard. 2003. "Consumption Risk and Cross-Sectional Returns." Working Paper no. 9538 (March), NBER, Cambridge, MA.
- Piazzesi, Monika. 2002. "Comment" [on Gabaix and Laibson, "The 6D Bias and the Equity-Premium Puzzle"]. In *NBER Macroeconomics Annual 2001*, vol. 16, edited by Ben S. Bernanke and Ken Rogoff. Cambridge, MA: MIT Press.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2003. "Housing, Consumption, and Asset Pricing." Manuscript, Dept. Econ., Univ. California, Los Angeles.
- Restoy, Fernando, and Philippe Weil. 1998. "Approximate Equilibrium Asset Prices." Working Paper no. 6611 (June), NBER, Cambridge, MA.
- Rubinstein, Mark. 1976. "The Valuation of Uncertain Income Streams and the Pricing of Options." *Bell J. Econ.* 7 (Autumn): 407–25.
- Startz, Richard. 1989. "The Stochastic Behavior of Durable and Nondurable Consumption." *Rev. Econ. and Statis.* 71 (May): 356–63.
- Wilcox, David W. 1992. "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work." *A.E.R.* 82 (September): 922–41.
- Yogo, Motohiro. 2003. "A Consumption-Based Explanation of Expected Stock Returns." Manuscript, Dept. Econ., Harvard Univ.