Idiosyncratic Consumption Risk and the Cross Section of Asset Returns
Author(s): Kris Jacobs and Kevin Q. Wang
Source: *The Journal of Finance*, Vol. 59, No. 5 (Oct., 2004), pp. 2211-2252
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/3694822

THE JOURNAL OF FINANCE • VOL. LIX, NO. 5 • OCTOBER 2004

Idiosyncratic Consumption Risk and the Cross Section of Asset Returns

KRIS JACOBS and KEVIN Q. WANG*

ABSTRACT

This paper investigates the importance of idiosyncratic consumption risk for the cross-sectional variation in asset returns. We find that besides the rate of aggregate consumption growth, the cross-sectional variance of consumption growth is also a priced factor. This suggests that consumers are not fully insured against idiosyncratic consumption risk, and that asset returns reflect their attempts to reduce their exposure to this risk. The resulting two-factor consumption-based asset pricing model significantly outperforms the CAPM, and its performance compares favorably with that of the Fama–French three-factor model.

THE VALIDITY OF THE WORKHORSE CONSUMPTION-BASED representative-agent model rests upon a number of assumptions, including the presence of markets that allow agents to perfectly insure themselves against idiosyncratic consumption risk. The only relevant pricing factor in the resulting standard multiperiod asset pricing model without frictions is the growth rate of aggregate consumption. The consumption-based asset pricing model (CCAPM) that reflects this allocation, however, is not able to match important aspects of the distribution of historical asset returns, such as the risk premium on the market portfolio. Its performance in a cross-sectional context is weak; as a result this model has not succeeded in replacing alternative cross-sectional models such as the capital asset pricing model (CAPM) as the standard in the profession.[1]

However, factors other than aggregate consumption growth become relevant to price assets if the assumptions underlying the standard representative agent model do not hold. For example, if agents cannot perfectly insure

*Jacobs is from the Faculty of Management, McGill University, and CIRANO; Wang is from the Joseph L. Rotman School of Management, University of Toronto. We would like to thank Joao Cocco, Francisco Gomes, Rick Green, Raymond Kan, Yuming Li, Tom McCurdy, Valery Polkovnichenko, Sergei Sarkissian, Raman Uppal, Petr Zemcik, an anonymous referee and seminar participants at the 2003 AFA meeting, the 2002 EFA meeting, the universities of Tilburg, Leuven, and Gent, Charles University in Prague, and the London Business School for helpful comments. We are also grateful to Eugene Fama and Raymond Kan for providing us with the asset return data. Jacobs acknowledges FCAR and SSHRC, and Wang acknowledges SSHRC and the Connaught Fund at the University of Toronto for financial support.

[1] A recent paper by Lettau and Ludvigson (2002) demonstrates that consumption-based models can challenge the CAPM along certain dimensions. This research is discussed below.

themselves against idiosyncratic consumption risk,[2] all higher moments of the cross-sectional distribution of consumption growth are relevant pricing factors. Researchers have long realized that changes in these moments may be of critical importance to explain changes in asset prices (see Mehra and Prescott (1985)). Building on this insight, a number of studies have investigated the importance of market incompleteness for the equity premium puzzle and the risk-free rate puzzle (see Telmer (1993), Constantinides and Duffie (1996), Heaton and Lucas (1996), Jacobs (1999), Vissing-Jorgensen (2002), Cogley (2002), Brav, Constantinides, and Geczy (2002), and Balduzzi and Yao (2000)). These studies provide mixed evidence on market incompleteness and the literature has not yet fully matured, but it appears that models with uninsurable idiosyncratic consumption risk and potentially limited market participation stand a better chance to explain the data than standard representative-agent models.

We investigate the importance of uninsurable idiosyncratic risk for the cross section of asset returns by testing if higher moments of the distribution of consumption growth are priced. Our econometric approach is designed to keep the empirical analysis relatively simple and to allow us to conduct a search over different specifications. We investigate a variety of pricing kernels that are linear in consumption growth and consumption dispersion. By focusing on the first two moments, we avoid the effects of possible mismeasurement of higher moments. By specifying a linear kernel, we greatly facilitate the econometric problem. This linear kernel can be viewed as an approximation to a utility-based nonlinear pricing kernel. We investigate the empirical performance of these pricing kernels using household consumption data from the Consumer Expenditure Survey (CEX). The main focus of our analysis is on kernels constructed using data on the consumption of nondurables and services for households that hold assets. One set of pricing factors is constructed using data on individual household consumption. Other pricing factors are obtained using data on the consumption of synthetic cohorts to address the presence of measurement error in household consumption data. Finally, whereas the time dimension of our empirical investigation is limited by the availability of the CEX data (1984–1995), we also use mimicking portfolio techniques that allow us to investigate the robustness of our results over a longer estimation period (1963–1995).

Our main conclusion is that both consumption growth and consumption dispersion are priced factors. However, the empirical findings confirm the importance of the measurement error problem. While estimation results for the consumption dispersion factor based on individual data are often insignificant

[2] At this point, it is important to elaborate on the terminology used in this paper in order to avoid confusion. In the literature on the CAPM, standard terminology splits up the risk of an individual asset into market risk and idiosyncratic risk. In this paper the focus is on idiosyncratic risk for an individual consumer. It is standard in the incomplete markets literature to refer to this risk as "idiosyncratic income risk" or "idiosyncratic risk." In this paper we do not investigate a full general equilibrium model but focus exclusively on equilibrium intertemporal consumption allocations. Therefore, we refer to this idiosyncratic risk as "idiosyncratic consumption risk." This terminology is imperfect but preferable to the use of "idiosyncratic risk," which could be confused with the terminology used in the context of the CAPM.

and not robust, our empirical results for the cohort-based pricing factors confirm the importance of uninsurable consumption risk for the cross section of asset returns. The results also indicate that the return on an asset is higher the more negatively correlated it is with consumption dispersion. This finding is of interest for the extensive literature on the equity premium puzzle, because one potential explanation for the puzzle is that consumption dispersion is larger in recessions. Intuitively, this leads to an increase in the risk faced by an individual agent, and this leads to a larger risk premium to induce investors to hold risky assets. Our estimates provide some support for this explanation.

To evaluate the significance of these findings, we investigate their robustness and compare the performance of the pricing factors against a number of alternatives. We find that the two-factor consumption-based model (with consumption growth and consumption dispersion as factors) outperforms the CAPM and the one-factor consumption-based model. Moreover, the empirical performance of the two-factor consumption-based model also compares favorably with that of the three-factor model proposed by Fama and French (1993). Finally, we investigate pricing kernels that combine size and book-to-market factors and/or the CAPM factor with the consumption-based factors. It is shown that even after accounting for these alternative pricing factors, the consumption-based factors are estimated significantly in the pricing equation. Moreover, statistical tests do not support dropping the consumption-based factors from a pricing kernel that also includes the Fama–French factors.

The paper proceeds as follows. Section I discusses related literature and motivates the empirical analysis. Section II describes the data and Section III presents the empirical results. Section IV concludes.

I. Idiosyncratic Consumption Risk and the Cross Section of Asset Returns

Following the seminal contributions by Lucas (1978) and Breeden (1979), a number of papers have conducted empirical investigations of representative-agent consumption-based asset pricing models. Even though these models have a wide range of empirical implications, a large part of the literature has a rather limited focus. In fact, much of the empirical research on consumption-based models has focused exclusively on the returns on a riskless asset and the market index, leading to the so-called equity premium and risk-free rate puzzles.[3] A small number of papers study the performance of the consumption-based model in a cross-sectional context. Mankiw and Shapiro (1986), Breeden, Gibbons, and Litzenberger (1989), and Cochrane (1996) conclude that the performance of the consumption-based model is unsatisfactory and that the consumption-based model performs no better than the CAPM. However, more recently Lettau and

[3] Hansen and Singleton (1982, 1984), Mehra and Prescott (1985), and Grossman, Melino, and Shiller (1987) focus exclusively on a riskless and a risky asset. Other papers such as Hansen and Singleton (1983) and Epstein and Zin (1991) focus on the equity premium puzzle but investigate some other risky assets. However, none of these papers focuses specifically on the cross section of returns.

Ludvigson (2002) show that these negative conclusions about the performance of the consumption-based model are due to the fact that these studies investigate an unconditional linear factor model. The performance of the conditional version of the model is about as good as that of the Fama–French three-factor model, when using a specific conditioning variable that is suggested by theory. Campbell and Cochrane (2000) provide an explanation for why consumption-based asset pricing models perform better conditionally than unconditionally.

This paper reaffirms that consumption-based asset pricing models are valuable for the study of the cross section of asset returns even when studying unconditional models, as opposed to the conditional models studied by Lettau and Ludvigson (2002). The key to this finding is that one has to move away from the rigid construction of a representative agent economy, which implies the irrelevance of idiosyncratic consumption risk. To appreciate the importance of this modeling approach, it is instructive to review the importance of complete markets and the representative agent assumption for the equity premium and risk-free rate puzzles.[4] The complete markets assumption is one of the assumptions underlying the representative agent model. Individual agents that are faced with a complete markets structure can insure themselves against idiosyncratic consumption risk. As a consequence, the prices of assets in the economy are equivalent to the prices in a closely related representative agent economy.

Whereas the complete markets assumption is a convenient modeling technique, casual observation as well as empirical testing has convinced most researchers that it is not very realistic (see Cochrane (1991), Mace (1991), and Hayashi, Altonji, and Kotlikoff (1994)). It is therefore not surprising that a growing number of studies investigate to what extent market incompleteness is of interest to explain the empirical rejections of the consumption-based models. A number of these studies investigate this issue by using simulation-based models. Whereas early studies by Telmer (1993) and Heaton and Lucas (1996) do not manage to generate large enough risk premia for most realistic parameterizations of the economy, later studies by Telmer, Storesletten, and Yaron (1997) and Constantinides, Donaldson, and Mehra (2002) have managed to generate larger risk premia under the assumption that idiosyncratic shocks are fairly persistent. A number of other studies (Jacobs (1999), Vissing-Jorgensen (2002), Cogley (2002), Brav, Constantinides, and Geczy (2002), and Balduzzi and Yao (2000)) have analyzed market incompleteness from another perspective, by investigating Euler equations that hold even if markets are incomplete.

[4] The literature contains other attempts to explain the equity premium puzzle and risk-free rate puzzle. A number of papers have focused on the importance of time aggregation (see Grossman, Melino, and Shiller (1987) and Heaton (1993)). Also, an extensive literature has studied the modeling of alternative preferences for the representative agent (see Abel (1990), Campbell and Cochrane (1999), Cochrane and Hansen (1992), Constantinides (1990), Detemple and Zapatero (1991), Epstein and Zin (1991), Ferson and Constantinides (1991), Heaton (1995), and Sundaresan (1989)). These approaches alleviate some of the problems with representative agent models and it is possible that they would also improve the cross-sectional performance of consumption-based models. See Kocherlakota (1996), Campbell, Lo, and MacKinlay (1997), and Cochrane (2001) for overviews of this literature.

Sarkissian (2003) analyzes incomplete risk sharing between countries. The test results in these papers are mixed, but a robust conclusion is that the degree of risk aversion implied by restrictions from incomplete markets is lower than that implied by representative agent models. Taken together, the findings in the literature on market incompleteness seem to indicate that accounting for idiosyncratic consumption risk has at least some potential to explain the structure of asset returns (see also Constantinides (2002)).

As models with incomplete markets have had some success explaining the equity premium and risk-free rate puzzles, it seems natural to investigate if they can be used to explain a wider cross section of asset returns. In cross-sectional asset pricing, the CAPM is the dominant paradigm. It is therefore a natural benchmark to evaluate the performance of a consumption-based model with market incompleteness. Ferson (1995) and Cochrane (1996, 2001) show that traditional factor pricing models such as the CAPM can be implemented by using the intertemporal optimality condition:

$$\mathrm{E}[M_t R_{j,t} \mid \Omega_{t-1}] = 1, \tag{1}$$

where M_t is the pricing kernel, $R_{j,t}$ is the return on asset j at time t, and Ω_{t-1} is the information set available to the econometrician at time $t-1$. For the CAPM, the pricing kernel M_t is specified as follows:

$$M_t = \beta_0 + \beta_1 R_{M,t}, \tag{2}$$

where $R_{M,t}$ is the return on the market portfolio at time t. We now outline a framework that allows us to compare the performance of a pricing kernel that accounts for idiosyncratic consumption risk with the performance of the CAPM as expressed in (1) and (2). The intertemporal optimality condition associated with individual i's investment in asset j implies that

$$\mathrm{E}[M_{i,t}(cg_{i,t}) R_{j,t} \mid \Omega_{t-1}] = 1, \tag{3}$$

where the pricing kernel $M_{i,t}$ that is indexed by individual i depends on consumption growth $cg_{i,t} = c_{i,t}/c_{i,t-1}$ in the context of a consumption-based asset pricing model. Averaging this orthogonality condition for asset j over all N consumers, we get

$$\mathrm{E}\left[\frac{1}{N}\sum_{i=1}^{N} M_{i,t}(cg_{i,t}) R_{j,t} \mid \Omega_{t-1}\right] = 1. \tag{4}$$

The pricing kernel in (4) will also be referred to as $M_t = (1/N)\sum_{i=1}^{N} M_{i,t}(cg_{i,t})$. Evaluating the performance of this kernel in the cross section can then be accomplished by specifying the underlying structure of the economy. For instance, if individual consumers have time-separable constant relative risk aversion (TS-CRRA), this average intertemporal Euler equation for asset j is

$$\mathrm{E}\left[\frac{1}{N} e^{-\theta} \sum_{i=1}^{N} (cg_{i,t})^{-\alpha} R_{j,t} \mid \Omega_{t-1}\right] = 1, \tag{5}$$

where α is the rate of relative risk aversion and θ is the rate of time preference.

Constantinides and Duffie (1996, henceforth CD) emphasize the importance of consumption dispersion under market incompleteness. Using a TS-CRRA specification, they specify an economy that leads to an Euler equation that specifies explicitly how the pricing kernel depends on the moments of the cross-sectional distribution of consumption growth. Specifically, their economy yields the following intertemporal Euler equation,

$$\mathrm{E}\left[e^{-\theta}\left(\frac{c_t}{c_{t-1}}\right)^{-\alpha}\exp\left(\frac{\alpha(\alpha+1)}{2}y_t^2\right)R_{j,t}\mid\Omega_{t-1}\right]=1, \tag{6}$$

where c_t is aggregate consumption at time t and y_t^2 can be interpreted as the variance of the cross-sectional distribution of $\log[(c_{i,t}/c_t)/(c_{i,t-1}/c_{t-1})]$.[5] In general, however, the pricing kernel will depend on all moments of the cross-sectional distribution, and it can be shown that (6) obtains as a special case (Mankiw (1986), Brav, Constantinides, and Geczy (2002), and Cogley (2002)).

The cross section of asset returns can be analyzed by using the generalized method of moments to evaluate (5) and (6) directly. However, the disadvantage of this approach is that the resulting econometric problem is highly nonlinear. This may complicate the optimization and the comparison with the benchmark CAPM because of the existence of local optima. We therefore use a different approach inspired by Cochrane (1996). With incomplete markets the pricing kernel depends on the cross-sectional moments of consumption growth.[6] Moreover, because it is difficult to estimate higher moments precisely, it may be preferable to limit attention to the first two moments. The precise nature of the relationship between the pricing kernel and these moments depends on the specification of the utility function. We therefore assume that the pricing kernel depends in a simple linear way on the first two cross-sectional moments of consumption growth:[7]

$$M_t = \beta_0 + \beta_1 mcg_t + \beta_2 vcg_t, \tag{7}$$

where $mcg_t = (1/N)\sum_{i=1}^{N}(cg_{i,t})$ and $vcg_t = (1/N)\sum_{i=1}^{N}(cg_{i,t} - mcg_t)^2$. This linear kernel (7) can be seen as a Taylor series approximation to an arbitrary

[5] Balduzzi and Yao (2000) use a different setup that leads to a different Euler equation. In their economy, the second factor is not the cross-sectional variance of log consumption growth, but the difference of the variance in cross-sectional consumption. They also use this kernel to study the cross section of asset returns.

[6] It must be noted that market incompleteness is not a necessary condition for the higher moments of consumption growth to enter the pricing kernel. For example, investor heterogeneity can have similar effects (Dumas (1989)). Basak and Cuoco (1998) explicitly characterize the cross-sectional distribution of consumption in a model with market incompleteness and investor heterogeneity.

[7] Heaton and Lucas (2000) also investigate linear pricing kernels with measures of idiosyncratic income risk as pricing factors. They find that the existence of entrepreneurial income risk has a significant influence on asset returns.

utility function. Appendix A discusses this in more detail and provides an example using TS-CRRA utility.

We emphasize that equation (7) indicates the relevance of consumption dispersion for the cross section of asset returns, as is also evidenced by the intertemporal optimality conditions (1), (3), (4), (5), and (6), which hold for every asset j. To see this more directly, consider rewriting (1) as (suppressing the dependence on the information set Ω_{t-1} for notational convenience)

$$\mathrm{E}(R_{j,t}) = \frac{1}{\mathrm{E}(M_t)} - \frac{\mathrm{cov}(M_t, R_{j,t})}{\mathrm{E}(M_t)}. \tag{8}$$

Substituting the pricing kernel (7) in (8) we get

$$\mathrm{E}(R_{j,t}) = \frac{1}{\mathrm{E}(M_t)} - \frac{\beta_1 \mathrm{cov}(mcg_t, R_{j,t})}{\mathrm{E}(M_t)} - \frac{\beta_2 \mathrm{cov}(vcg_t, R_{j,t})}{\mathrm{E}(M_t)}, \tag{9}$$

which clearly indicates that consumption dispersion affects the return for every asset j, and therefore the cross section of asset returns.

Expression (9) illustrates a well-known caveat with the interpretation of the coefficients in the pricing kernel (7). Consider the coefficient β_1 on mean consumption growth. We expect that an asset that is more highly correlated with consumption growth pays a higher return, because it is less attractive. From (9), it is clear that this positive correlation requires a negative estimate for β_1. Using (9) and Appendix A, we can also provide some intuition about the sign of the coefficient β_2 on consumption dispersion. Appendix A indicates that the sign of β_2 depends on the second derivative of the pricing kernel, or equivalently, the third derivative of the utility function. For a TS-CRRA utility function, this leads us to expect a positive β_2. In other words, we expect that an asset that pays a high return when the variance of consumption growth is high has a low expected return, because it is a relatively more valuable asset. Although other often-used utility functions yield a similar result, Mankiw (1986) correctly argues that a negative sign is also possible.[8]

Finally, we also compare the performance of the consumption-based factors to a benchmark other than the CAPM. We use the size and book-to-market factors proposed by Fama and French (1992, 1993), and we use the kernel for the Fama–French three-factor model

$$M_t = \beta_0 + \beta_1 R_{M,t} + \beta_2 \mathrm{SMB}_t + \beta_3 \mathrm{HML}_t, \tag{10}$$

where SMB_t is the size factor and HML_t is the book-to-market factor. To evaluate the relative performance of the consumption-based factors compared to the Fama–French factors, we investigate a number of kernels that interact the

[8] The literature on the equity premium puzzle has attempted to explain the historical equity premium with a combination of a positive β_2 and a negative covariance between the time series of consumption dispersion and the time series of equity returns. Intuitively, larger consumption dispersion in periods of low equity returns (recessions) leads to an increase in the risk faced by an individual agent, and a larger risk premium to induce investors to hold risky assets (see CD (1996) and Cogley (2002)).

consumption-based factors with the size and/or book-to-market and/or market factors. These kernels are described in more detail in the tables. Using the Fama–French factors as a benchmark is a logical choice because they have achieved some degree of empirical success. It must be noted, however, that these factors are essentially a function of prices and do not address the underlying relationship with economic quantities. Lettau and Ludvigson (2002) use the consumption–wealth ratio, but wealth is also price-related. In contrast, our pricing kernel (7) attempts to explain the cross section of returns using economic quantities, just like the pricing kernel for a representative-agent consumption-based model.

II. Data

This section briefly discusses the construction of the consumption and asset return data. Further details on data construction are presented in Appendix B. The empirical tests are implemented as follows. First, consumption data are used to construct pricing factors that are the first and the second moments of the cross-sectional distribution of consumption growth. In the second stage, these pricing factors are taken as given in an econometric investigation of the intertemporal relation (1) for a cross section of asset returns. It must therefore be noted that the uncertainty involved in constructing the pricing factors is neglected in this econometric analysis. The consumption data from the CEX that are used in most of the empirical tests are discussed in Section II.A. The cross section of asset returns is described in Section II.B. Section II.C motivates the use of synthetic cohorts. Section II.D discusses the Consumer Population Survey (CPS), which is used as an alternative to the CEX in some of the empirical tests. Finally, Section II.E provides a detailed discussion of the statistical properties of the data used in the empirical analysis and the factors used in the pricing equation.

A. Consumer Expenditure Survey (CEX) Data

To construct the pricing factors used in most of the empirical work, we use data on household consumption from the CEX. The CEX data have been used by a number of researchers to analyze the importance of idiosyncratic consumption risk for the equity premium puzzle.[9] We construct the pricing factors using a measure of monthly nondurable plus services (NDS) consumption extracted from the CEX for the period April 1984 to September 1995. Appendix B describes the construction of this consumption measure in detail. Most of our empirical analysis is done using a sample of NDS consumption that only contains asset holders. The definition of asset holders is also given in Appendix B.

[9] See Brav, Constantinides, and Geczy (2002), Vissing-Jorgensen (2002), Cogley (2002), and Balduzzi and Yao (2000). Balduzzi and Yao (2000) also present an analysis of cross-sectional pricing using their (different) pricing kernel.

We repeated our analysis for a sample that includes all consumers, and we report some of those results in Appendix C. It must be noted though that the interpretation of those results may be problematic, because the Euler equation and the resulting pricing relationship do not hold for consumers who are not at an interior solution.

B. Asset Return Data

We use a set of test portfolios that includes the 25 size and book-to-market portfolios of Fama and French (1993), a long-term government bond, a long-term corporate bond, and the 3-month Treasury bill rate. The construction of these asset returns is described in some detail in Appendix B. The empirical tests are run twice: Once with the 25 equity portfolios and a second time with all 28 assets. It is worthwhile to inspect the descriptive statistics for the test portfolios for the sample period April 1984 to September 1995 in Panel A of Table I. An important observation is that for the stock portfolios, which are sorted according to size and book-to-market, the pattern of average returns

Table I

Summary Statistics for Asset Returns

This table presents means and standard deviations for the test asset returns for the sample periods 1984–1995 (Panel A) and 1963–1995 (Panel B). The means and standard deviations of real monthly returns on the 25 size and book-to-market portfolios of Fama and French (1993) are first presented. The notations SZ1 through SZ5 stand for the five size quintiles (from small to large), while BM1 through BM5 stand for the five book-to-market equity quintiles (from low to high). Next, the real returns on government bonds (GBT: CRSP variable GBTRET), corporate bonds (CBT: CRSP variable CBTRET), and Treasury bills (TB) are included.

	Means					Standard Deviations				
	BM1	BM2	BM3	BM4	BM5	BM1	BM2	BM3	BM4	BM5
					Panel A: 1984–1995					
SZ1	−0.007	0.018	0.018	0.027	0.029	0.104	0.092	0.086	0.078	0.091
SZ2	0.012	0.023	0.032	0.031	0.030	0.094	0.088	0.075	0.070	0.091
SZ3	0.020	0.027	0.023	0.031	0.035	0.089	0.080	0.070	0.069	0.076
SZ4	0.026	0.024	0.024	0.029	0.033	0.076	0.069	0.076	0.062	0.070
SZ5	0.028	0.027	0.023	0.026	0.039	0.070	0.066	0.052	0.055	0.070
	GBT	CBT	TB			GBT	CBT	TB		
	0.022	0.021	0.006			0.046	0.038	0.005		
					Panel B: 1963–1995					
SZ1	0.014	0.025	0.027	0.033	0.037	0.126	0.111	0.103	0.096	0.103
SZ2	0.015	0.024	0.030	0.033	0.036	0.111	0.099	0.090	0.085	0.093
SZ3	0.017	0.025	0.025	0.029	0.033	0.101	0.089	0.079	0.077	0.086
SZ4	0.018	0.018	0.025	0.029	0.032	0.089	0.081	0.076	0.074	0.086
SZ5	0.015	0.016	0.017	0.020	0.027	0.073	0.067	0.058	0.060	0.072
	GBT	CBT	TB			GBT	CBT	TB		
	0.007	0.008	0.003			0.045	0.043	0.007		

is very different from the one documented in Fama and French (1993). This difference is due to the different sample period. Most importantly, for the sample under consideration in this paper, Panel A of Table I confirms the observation of Cochrane (2001, p. 438) that the size effect has disappeared in the 1980s. It will be argued below that this data problem, due to the limited time dimension of the consumption data and the unusual pattern of the returns data over this period, will somewhat complicate the interpretation of our results. We therefore also report test results based on the sample period July 1963 to December 1995. These tests are based on a mimicking portfolio approach that is described below. Panel B of Table I presents descriptive statistics for these returns data. It can be seen that in contrast with the returns in Panel A, the returns in Panel B display important size (SMB) and book-to-market (HML) effects. This is not surprising, because the sample period in Panel B is very similar to the one in Fama and French (1993).

C. Constructing Synthetic Cohorts

We start by constructing pricing factors using the cross section of individual consumption growth at every time t. Subsequently, we use these pricing factors in a cross-sectional pricing relationship. The problem with this approach is the existence of measurement error in household consumption data, which is well documented. To deal with the measurement error problem, we adopt the synthetic cohorts approach that is popular in the economics literature.[10] This approach basically involves the construction of a representative consumer for a typical group that can be defined by observable characteristics. Most of our empirical results are obtained using synthetic cohorts based on age and education. We also provide a robustness analysis using cohorts based exclusively on age. In our empirical analysis, we refer to the factor for mean consumption growth based on individual consumption data as $mcg_{i,t}$ and to the corresponding variance factor as $vcg_{i,t}$. The mean and variance factors based on the age-and-education cohorts are denoted as $mcg_{ae,t}$ and $vcg_{ae,t}$, while the mean and variance factors based on the age cohorts are denoted as $mcg_{a,t}$ and $vcg_{a,t}$. The construction of the cohorts is discussed in more detail in Appendix B.

D. Consumer Population Survey (CPS) Data

While we are not aware of other U.S. data sets that contain detailed information on consumption growth going back further in time than the CEX,[11] there are some alternative sources of information that can be used to further investigate the importance of consumption dispersion. Several data sets contain high-quality information on income data, and to the extent that income

[10] The measurement error problem is documented in (among others) Altonji (1986), Altonji and Siow (1987), and Zeldes (1989). The synthetic cohort approach was previously used by (among others) Browning, Deaton, and Irish (1985) and for the CEX data by Attanasio and Weber (1995).

[11] The Panel Study of Income Dynamics (PSID) contains consumption data at a yearly frequency going back to 1970, but only data on food consumption are available.

growth is correlated with consumption growth, we can use factors constructed using income growth to proxy for consumption-based factors. We therefore repeat our empirical tests using income data from the CPS. The CPS contains cross-sectional income data for the years 1966–2000. We use cohorts from the CPS by ranking households according to income, using 20 income groups in each cross section. Subsequently, we construct income growth data by considering income growth in a given year for the average person in the poorest cohort, the second-poorest cohort, and so forth. In this way, we obtain a time series of income growth and income dispersion using 20 cross-sectional cohorts for the years 1967–2000.[12]

E. Descriptive Statistics

Table II presents descriptive statistics on individual and cohort consumption growth. The following observations are of interest. First, the distribution of individual consumption growth is not adequately described by the normal distribution, with the statistics indicating positive skewness and excess kurtosis.[13] Not surprisingly, the distribution of cohort consumption growth is more adequately described by a normal distribution. While it is tempting to attribute these differences (especially the lower variance) to the elimination of measurement error, it is also possible that by constructing the cohorts, we have eliminated some genuine variability in consumption that is the result of unanticipated shocks that were not fully insured. Second, the mean consumption growth rate in Table II is much higher than the one obtained from NIPA data (not reported). The key to this finding is of course that NIPA growth rates are obtained by averaging over consumption levels, and not over consumption growth as in Table II.[14] Third, the distribution of consumption growth for asset holders does not seem to differ dramatically from the distribution of consumption growth

[12] We are not aware of data sources that allow construction of longer time series on income growth and income dispersion, even though standard measures of income dispersion such as Gini coefficients are available for longer time periods. Such standard measures are not useful for our purposes: We need raw data on income growth to construct factors that are as close as possible to the theoretical construct of consumption growth that we are testing.

[13] The distribution of consumption growth is bounded below by zero. The distribution of the logarithm of consumption growth (which is used in the CD pricing kernel (6)) is not bounded from below, and therefore deviations from normality are less pronounced. Using factors based on the logarithm of consumption growth, however, complicates reporting in the case of the cohort-based factors, because it is not clear if one should first take logarithms and then construct cohorts or first construct cohorts and then take logarithms. It must also be noted that we construct the pricing factors for every period separately. Therefore, some of the deviations from normality evident in Table II are caused by fluctuations of consumption growth over time that do not show up in the pricing factors. We therefore also computed the first four moments of the data on a month-by-month basis (not reported). The resulting skewness and excess kurtosis coefficients are lower, but still indicate deviations from normality. Large outliers are also very rare in these data.

[14] We computed aggregate growth rates from the CEX to investigate the reliability of the data. While there are some interesting differences between the NIPA and the numbers constructed from the CEX, the average growth rate over the sample is very similar.

Table II

Summary Statistics for Consumption Growth

This table presents descriptive statistics for individual and cohort consumption growth and the level of household consumption. It presents the first four moments, the minimum, and the maximum. It also presents the average cohort size, as well as the minimum and maximum cohort sizes for the cohort-based statistics. The descriptive statistics are for the 1984–1995 sample consisting of 134 monthly observations. NDS stands for nondurable and services consumption, AH denotes asset holders, and ALL indicates that the sample includes all households.

	No. Obs	Mean	SD	Skew	Exc. Kurt	Min	Max
			Panel A: Individual Consumption Growth				
NDS, ALL	1,04,810	1.060	0.431	5.407	154.330	0.052	25.106
NDS, AH	33,855	1.053	0.411	5.887	167.212	0.104	20.228
			Panel B: Individual Consumption Level (in $ of 1984)				
NDS, ALL	1,04,810	2,166	1,326	3.099	30.187	100.56	43,318
NDS, AH	33,855	2,430	1,388	3.512	44.970	114.86	43,318

	No. Obs	Mean	SD	Skew	Exc. Kurt	Min	Max	Av.Cs	Min.Cs	Max.Cs
				Panel C: Cohort Consumption Growth, Age Cohorts						
NDS, ALL	1,608	0.999	0.075	0.307	0.101	0.859	1.360	65.18	28	121
NDS, AH	1,608	1.000	0.110	0.457	1.771	0.629	1.762	21.05	3	46
				Panel D: Cohort Consumption Level, Age Cohorts (in $ of 1984)						
NDS, ALL	1,608	2,172	352	0.044	−0.271	1248	3344	65.18	28	121
NDS, AH	1,608	2,413	449	0.330	0.494	1107	4561	21.05	3	46
				Panel E: Cohort Consumption Growth, Age-and-Education Cohorts						
NDS, ALL	1,608	0.997	0.083	0.157	0.809	0.627	1.376	65.18	10	165
NDS, AH	1,608	1.001	0.116	0.711	2.510	0.650	1.682	21.05	2	48
				Panel F: Cohort Consumption Level, Age-and-Education Cohorts (in $ of 1984)						
NDS, ALL	1,608	2,331	582	0.983	0.815	1268	4,776	65.18	10	165
NDS, AH	1,608	2,473	544	0.907	1.244	1152	4,911	21.05	2	48

based on all consumers. Fourth, Table II strongly suggests the presence of measurement error in these data. For instance, some households consume 20 times as much or 10 times less in a given 3-month period compared to the previous period. In one instance, a household only consumes 5% of its lagged consumption. However, these outliers are unusual. There are very few cases in which consumption increases more than fivefold. Table II also contains information on cohort construction. Some cohorts contain few observations, but on the positive side, the average cohort size is fairly large in all cases. Also, as expected, the average cohort size is much larger for the sample consisting of all consumers, as compared to the sample consisting of asset holders only.

Table III presents the descriptive statistics for the pricing factors. It suggests a number of stylized facts that will be helpful in interpreting the test results.

Table III

Summary Statistics for Consumption-Based Pricing Factors

This table presents descriptive statistics for the consumption-based pricing factors. It presents the mean, maximum, minimum, and standard deviation of each factor in the 1984–1995 sample, which consists of 134 monthly observations. One set of pricing factors is based on consumption data for individual households

$$mcg_{i,t} = \frac{1}{N}\sum_{i=1}^{N}(cg_{i,t}), \quad vcg_{i,t} = \frac{1}{N}\sum_{i=1}^{N}(cg_{i,t} - mcg_{i,t})^2,$$

where $cg_{i,t} = (c_{i,t}/c_{i,t-1}), c_{i,t}$ is the consumption of individual i at time t, and N is the number of households in the cross section. For the cohort-based factors, consider the average consumption of a representative cohort k at time t, that is, $\text{cohc}_{k,t} = (1/N_{k,t})\sum_{i=1}^{N_{k,t}}(c_{i,t})$, where $N_{k,t}$ is the number of observations on this cohort at time t. Define $\text{coh}cg_{k,t} = (\text{cohc}_{k,t}/\text{cohc}_{k,t-1})$. Denoting the cohort construction based on age with the subscript a, the factors are

$$mcg_{a,t} = \frac{1}{H}\sum_{k=1}^{H}(\text{coh}cg_{k,t}), \quad vcg_{a,t} = \frac{1}{H}\sum_{k=1}^{H}(\text{coh}cg_{k,t} - mcg_{a,t})^2,$$

where $\text{coh}cg_{k,t}$ is the growth rate of the average consumption for the age cohort k. The subscript *ae* stands for the cohort construction based on age and education.

	mcg_i	vcg_i	mcg_a	vcg_a	mcg_{ae}	vcg_{ae}
			Panel A: Asset Holders			
Mean	1.0535	0.1663	1.0002	0.0098	1.0011	0.0109
Min	0.9303	0.0750	0.8545	0.0020	0.8642	0.0017
Max	1.1454	1.4500	1.1026	0.0602	1.1159	0.0462
SD	0.0540	0.1268	0.0565	0.0067	0.0594	0.0083
			Panel B: All Households			
Mean	1.0602	0.1838	0.9969	0.0034	0.9975	0.0045
Min	0.9583	0.1010	0.8937	0.0004	0.8890	0.0001
Max	1.1641	0.8956	1.1002	0.0089	1.1040	0.0178
SD	0.0495	0.0820	0.0509	0.0016	0.0527	0.0027

First, while the differences between *mcg* factors based on individual data and cohorts are fairly modest, *vcg* factors for individual data and cohorts are often quite different. It will therefore probably be the case that the adoption of the cohort construction influences the estimation of the sign and magnitude of the second factor much more than that of the first factor. Second, there are some differences between the pricing factors constructed using asset holders only and those constructed using data on all households.

III. Empirical Findings

A. Testing Methods

To evaluate the significance of the cross-sectional dispersion of consumption growth, we use the generalized method of moments (GMM, Hansen (1982)).

This testing method has recently been implemented in various empirical studies of cross-sectional asset pricing. For example, see Cochrane (1996), Jagannathan and Wang (1996), and Heaton and Lucas (2000). Cochrane (2001, Chapter 15) demonstrates that the GMM approach works well for linear asset pricing models. We test the unconditional version of the orthogonality conditions

$$\mathrm{E}[M_t(\beta)R_{j,t}] = 1, \tag{11}$$

where $R_{j,t}$ is the return on the j^{th} test asset and $M_t(\beta)$ is the pricing kernel. We provide test results for the kernels discussed in Section I and some combinations of the pricing factors discussed there. Specifically, we consider pricing kernels of the form

$$\sum_k b_k f_{k,t} = b'f_t, \tag{12}$$

where f_t is the vector of factors and b is a vector of constant parameters. It is now well known that the above linear pricing kernel represents a multifactor model, which can be equivalently expressed in a linear multifactor beta pricing form (see Ferson (1995) and Cochrane (1996) for details).

One part of our testing strategy uses a standard iterated GMM testing procedure. After an initial round, we go through an iterating procedure where the weighting matrix is given by the sample covariance matrix of the orthogonality conditions evaluated at the estimate of β obtained in the previous round. This procedure is repeated until the estimates converge. Using the iterated estimates, we then compute Hansen's J-test statistic of overidentifying restrictions. Because the iterated GMM estimates are asymptotically statistically efficient, they form a natural starting point for investigating the statistical significance of the pricing factors.[15]

While the iterated GMM procedure yields efficient estimates, it implies that different assets in the sample are weighted on the basis of statistical efficiency. The resulting estimates may therefore be hard to interpret from an economic perspective (see Cochrane (1996, 2001)). We check the robustness of our results by also examining first- and second-stage GMM results. The first-stage GMM results are of great interest as a robustness check because the weighting of the different assets does not depend on statistical precision in this case. In principle, one can choose different weighting matrices for these first-stage estimates. We estimate parameter values that minimize the Hansen and Jagannathan (1991, 1997) distance measure of pricing errors.[16] This measure is computed as follows. Let

$$\mu_T = \frac{1}{T}\sum_{t=1}^{T}[M_t(\beta)R_t - 1], \tag{13}$$

[15] Two-stage GMM and iterated GMM are both asymptotically efficient. Ferson and Foerster (1994) use an extensive Monte Carlo analysis to demonstrate that the iterated procedure is preferable to the two-stage GMM procedure in realistic finite sample settings.

[16] These first-stage estimates are also used as the starting point for the iterated GMM procedure.

where R_t is the vector of returns on the test assets. The weighting matrix for the HJ distance (1997) is

$$W_{HJ} = \frac{1}{T}\sum_{t=1}^{T} R_t R_t'. \tag{14}$$

The HJ distance is then given by

$$d = (\mu_T' W_{HJ}^{-1} \mu_T)^{\frac{1}{2}}. \tag{15}$$

In other words, different orthogonality restrictions are weighted using data on returns only. The attractiveness of this weighting matrix is based upon the ability to interpret the resulting HJ measure as the maximum pricing error among all portfolio payoffs that have a unit second moment. It is also the least-square distance between the given candidate pricing kernel and the nearest point to it in the set of all pricing kernels that price assets correctly. Moreover, the measure is robust to portfolio formation. See Hansen and Jagannathan (1997) for details. As another robustness check, we also report adjusted R^2s for cross-sectional regressions that use the pricing factors under consideration. Whereas the HJ distances do not weigh different assets according to statistical precision, the weights in (14) can still involve large long and short positions. While it is harder to link the adjusted R^2s from cross-sectional OLS regressions to the intertemporal asset pricing framework underlying (11), they are easy to understand and can be used to compare models because they are based on equally weighted pricing errors.

B. Test Results

B.1. Test Results for the 1984–1995 Sample

Table IV presents estimation and test results for the April 1984 to September 1995 sample, obtained using the iterated GMM procedure for the sample that includes 25 equity portfolios. The consumption factors for this data set are directly computed from the CEX data. Each row represents a set of estimation results for (11) and the J-statistic associated with the estimation exercise is listed in the next to last column. The HJ distance is obtained from the first iteration of the corresponding GMM problem.

Consider the results associated with the pricing factors based on the individual consumption data in row 5. The consumption growth factor *mcg* is estimated with a negative sign and the *vcg* factor is estimated with a positive sign, but insignificantly. Now consider the estimation results obtained using cohort-based factors in row 6. The *mcg* factor is estimated with a negative sign and the *vcg* factor has a positive sign, and both are statistically significant.[17] The inclusion

[17] Note that positive correlation (conditional on the other factors) between the asset returns and the *mcg* factors shows up with a negative sign, and negative correlation (conditional on the other factors) between the asset returns and the *vcg* factors shows up with a positive sign.

Table IV

Testing for the Significance of Consumption-Based Factors

This table includes results for 15 linear pricing kernels, which are various combinations of the consumption-based factors, the cross-sectional mean and variance of consumption growth (mcg_t and vcg_t), and the Fama–French three factors, the market factor ($R_{M,t}$), and the size and book-to-market factors (SMB_t and HML_t). The sample period is 1984–1995. The test assets consist of the 25 Fama–French size and book-to-market portfolios. A standard GMM procedure is implemented for testing the moment conditions $\mathrm{E}[M_t(\beta)R_{it}] = 1$, where R_{it} is the return on the i^{th} test asset. In the initial round, the HJ distance measure is minimized. Then the iterated GMM estimates are obtained; that is, at each round, the weighting matrix is updated using the estimates from the previous round, and the procedure is repeated until estimates converge. Reported in the table are the iterated estimates and the J test statistics that are based on the iterated estimates. In parentheses under the estimates are t-statistics and in parentheses under the J statistics are the p-values. The last column includes the HJ distances based on the first-stage estimates. The tests are conducted using nondurable and services consumption for asset holders. The consumption factors are the two pairs based on consumption growth (mcg_i, vcg_i) and (mcg_{ae}, vcg_{ae}), defined in Table III (or see Section II and Appendix B).

Row	Const.	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	J-Test	HJ-Distance
1	28.80	−26.34							81.80	1.905
	(6.03)	(−5.88)							(0.000)	
2	21.91			−21.00					90.43	1.889
	(6.30)			(−6.14)					(0.000)	
3	24.71	−33.08				10.88	−7.64	−11.11	69.92	1.778
	(4.09)	(−5.54)				(2.49)	(−1.51)	(−2.54)	(0.000)	
4	16.52			−31.64		15.44	−12.61	−9.52	66.45	1.761
	(2.88)			(−5.56)		(3.40)	(−2.41)	(−2.20)	(0.000)	
5	30.64	−28.69	3.99						79.24	1.863
	(6.07)	(−5.84)	(1.01)						(0.000)	
6	16.09			−18.51	286.47				73.72	1.791
	(3.90)			(−4.29)	(5.20)				(0.000)	
7	24.47	−31.44	4.64			8.81			73.12	1.746
	(4.30)	(−5.69)	(1.07)			(2.65)			(0.000)	
8	7.67			−22.55	292.13	12.26			57.88	1.644
	(1.31)			(−4.28)	(5.06)	(3.76)			(0.000)	
9	31.61	−29.27	2.12				0.91	−9.23	77.32	1.819
	(6.03)	(−5.78)	(0.63)				(0.23)	(−2.28)	(0.000)	
10	29.30			−32.19	355.42		−2.44	−21.67	45.50	1.685
	(4.21)			(−4.51)	(4.74)		(−0.42)	(−4.47)	(0.001)	
11	26.10	−37.63	9.33			12.73	−7.11	−11.44	66.67	1.693
	(4.10)	(−5.37)	(1.33)			(2.50)	(−1.34)	(−2.59)	(0.000)	
12	15.20			−31.93	338.08	13.57	−19.76	−13.79	41.27	1.553
	(1.88)			(−4.56)	(4.76)	(2.80)	(−2.61)	(−2.69)	(0.002)	
13	−5.65					6.48			108.01	1.892
	(−3.86)					(4.48)			(0.000)	
14	1.00						2.34	−3.00	109.55	1.950
	(21.47)						(1.05)	(−1.68)	(0.000)	
15	−16.64					17.18	−20.20	−2.66	92.84	1.880
	(−5.67)					(5.99)	(−6.25)	(−1.22)	(0.000)	

of the market factor does not change inference on the consumption-based factors. However, the coefficient on the market factor is estimated with the wrong sign. The J-statistics for the different estimation exercises vary widely. It is tempting to interpret these differences, but the different rows do not represent nested hypotheses and therefore such an interpretation is not possible. We can only conclude that all test statistics indicate rejection of the null hypothesis at conventional levels of significance. However, the HJ distances indicate that the consumption-based models improve on the CAPM. Given that this is an unconditional model and that we know that one-factor consumption-based models do not perform well in an unconditional setting, we therefore implicitly conclude that the extra second factor adds explanatory power. Rows 1 and 2 verify this conclusion by presenting test results for pricing kernels including only the *mcg* factor. The results indicate that the performance of this consumption-based model is not very different from that of the CAPM, judging by the HJ distance measures. We can therefore safely conclude that it is the inclusion of the cross-sectional variance factor that drives the HJ distances down.

Rows 9–12 report on a more ambitious exercise designed to subject the consumption-based pricing factors to a potentially more stringent test. Instead of using the CAPM as a benchmark, we compare the performance of the consumption-based factors to the three-factor model proposed by Fama and French (1992, 1993), which includes size and book-to-market factors as well as the market portfolio. The results are interesting. First, when combining the consumption-based factors in a kernel with (a subset of) the three factors, the consumption-based factors show up significantly with the expected sign. Second, when adding the consumption-based factors to (subsets of) the Fama–French factors, the resulting HJ statistic is significantly lower. It is also interesting that the resulting HJ distances are usually lower than the ones obtained for the pricing kernel with two consumption-based factors only. While one has to keep in mind that there is no adjustment for the number of factors in the computation of the HJ distance, this finding may not necessarily be surprising given the fact that the size and book-to-market factors are so successful in capturing empirical patterns in stock returns (Fama and French (1996)). In other words, this result is probably as indicative of the explanatory power of the Fama–French factors as of the performance of the consumption-based factors. This conclusion is reinforced by the results in rows 3 and 4. When adding the Fama–French factors to a single consumption-based factor, the HJ distances decrease.[18]

There is one important caveat regarding the Fama–French factors. The point estimates and t-statistics in Table IV indicate that the HML factor and especially the SMB factor are often estimated insignificantly and sometimes with the wrong sign. These findings are to some extent foreshadowed by the descriptive statistics on the asset returns in Panel A of Table I. In the 1984–1995

[18] The fact that the HJ distances are all fairly high is due to the particular cross section under investigation, and does not mean that the models are inappropriate. Because we use a challenging cross section of portfolios sorted by size and book-to-market, it is to be expected that the HJ distance is high. The main focus is on the differences in HJ distance between different factor pricing models.

sample period, the book-to-market effect is not as prevalent as in the Fama–French (1993) sample, and the size effect is nonexisting. It is therefore perhaps not surprising that these factors are often estimated insignificantly. These findings indicate in a forceful way that our sample period is not only very short, but perhaps also not representative of historical properties of asset returns. This finding motivates some of the work below, which investigates the robustness of our findings using longer sample periods.

Table V is similar to Table IV, except that the returns used in estimation include the returns on two long-term bonds and the T-bill rate, as well as the 25 equity portfolios. These results can be summarized very briefly. First, the estimates of the coefficients on the *vcg* factor in Table V confirm the results in Table IV, and the HJ distances for the kernels that include the

Table V

Test Results Based on a Larger Set of Assets

This table presents results using the Fama–French 25 portfolios plus the two long-term bonds and the T-bill. It is otherwise identical to Table IV.

Row	Const.	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	J-Test	HJ-Distance
1	19.64	−17.59							112.65	2.751
	(4.97)	(−4.75)							(0.000)	
2	5.33			−4.34					118.99	2.759
	(2.08)			(−1.70)					(0.000)	
3	32.38	−20.87				−8.70	13.81	9.23	102.28	2.672
	(5.12)	(−4.07)				(−2.72)	(2.87)	(2.12)	(0.000)	
4	22.36			−9.59		−10.99	19.95	7.14	109.75	2.689
	(4.58)			(−2.45)		(−4.03)	(4.81)	(1.97)	(0.000)	
5	−4.99	14.53	−61.05						72.82	2.714
	(−0.62)	(1.64)	(−4.23)						(0.000)	
6	0.94			−3.72	324.72				91.63	2.696
	(0.29)			(−1.08)	(5.28)				(0.000)	
7	6.45	12.99	−61.60			−9.66			68.91	2.684
	(0.67)	(1.44)	(−4.34)			(−2.30)			(0.000)	
8	0.77			−3.83	327.17	0.25			91.30	2.689
	(0.19)			(−1.10)	(5.14)	(0.10)			(0.000)	
9	0.12	9.53	−59.23				−4.08	−13.74	69.23	2.686
	(0.01)	(1.01)	(−3.99)				(−0.60)	(−2.34)	(0.000)	
10	11.60			−15.53	462.15		−8.36	−20.17	64.21	2.652
	(2.20)			(−2.74)	(5.22)		(−1.23)	(−4.20)	(0.000)	
11	11.49	13.69	−61.38			−15.10	11.27	−12.44	63.75	2.620
	(1.15)	(1.44)	(−4.19)			(−2.99)	(1.37)	(−2.12)	(0.000)	
12	21.77			−16.17	409.14	−8.69	1.03	−25.85	69.82	2.610
	(2.56)			(−2.79)	(4.99)	(−1.79)	(0.14)	(−4.51)	(0.000)	
13	3.66					−2.63			118.74	2.743
	(2.19)					(−1.62)			(0.000)	
14	1.07						−3.23	−7.94	115.62	2.749
	(23.05)						(−1.45)	(−4.87)	(0.000)	
15	15.19					−13.53	20.31	−2.96	115.80	2.699
	(6.94)					(−6.47)	(6.64)	(−1.34)	(0.000)	

consumption-based factors are lower. Interestingly however, some of the estimates on the *mcg* factor do not have the expected sign, and in contrast with Table IV the estimates on the market factor do have the expected sign. It seems to be the case that the addition of the long-term bonds and the T-bill to the sample significantly affect the estimation results.

The tables in Appendix C provide some perspective on the robustness of these results. Table CI reports results similar to those in Table IV, but the cohorts are constructed using age. It can be seen that the results are fairly similar to those in Table IV. We have also analyzed the robustness of the results to changes in the composition of the age or education groups (not reported). The results of this robustness analysis can be summarized as follows. The estimates obtained using the cohort-based factors are robust; the results obtained using individual consumption data are not. This may suggest that estimation results are sensitive to measurement error in the individual data and that this measurement error is largely removed by the cohort approach. However, this statement has to be qualified somewhat: There are some instances of cohort-based results where the sign changes or the estimate is insignificant. Panel A of Table CIII further investigates the robustness of the results obtained in Table IV, by presenting the first-stage GMM estimates obtained by minimizing the HJ distance measure (15). We obtain positive estimates for the variance factors when we use factors based on cohort construction. It must be noted that in most cases the point estimates are not as significant as the ones in Table IV, but this is to be expected because the first-stage GMM estimation is less efficient. Panel B of Table CIII provides two-step GMM estimates. Again, the results confirm those of Table IV. The coefficients on the variance factors are estimated significantly positive whenever we use cohorts to construct pricing factors. It is also of interest that in most cases, the point estimates in Panels A and B of Table CIII are quite similar to the corresponding estimates in Table IV.

Finally, Table CII repeats the analysis of Table IV, but uses consumption data on all households, rather than asset holders only. As mentioned before, these results are not well motivated theoretically, but they are of some interest because the consumption patterns of asset holders and nonasset holders over the business cycle display some similarities, and the data set including nonasset holders is much larger. Interestingly, the results in Table CII largely confirm those of Table IV. This is almost invariably the case for other empirical results obtained using all consumers too, but they are omitted to keep the focus on the results for asset holders, which are motivated by theory. One interesting difference is that the point estimate for the *vcg* factors is always higher when all households are included in the sample. This reflects the higher variance of the dispersion factor when only asset holders are used in the construction of the factors (see Table III). We perform a number of other robustness exercises that are not presented in the tables because of space constraints. We repeat the analysis after deseasonalizing the data using the census X11 method implemented in EVIEWS. Second, we correct the household consumption data for family size in two different ways: By computing per capita data and by correcting for household size using a scale that is estimated from the data. Third, to address the problem that some cohorts contain few observations, we construct cohort pricing factors that are weighed by the size of the cohort. None of these adjustments cause significant changes in the empirical results.

B.2. Test Results for the 1963–1995 Sample

It was mentioned before that the statistical properties of returns in the 1984–1995 sample are not representative of other periods. It is therefore important to investigate the robustness of our results using longer samples, which is problematic as consumption data are not available. To cope with this problem, we use mimicking portfolio techniques to project the consumption factors onto longer time horizons. Suppose that consumption-based factors are obtained from a sample period of T_1 observations. Due to the data constraint, this is the maximum sample size available at the moment. However, the asset return data are available for a sample period of T_2 observations, where T_2 is much larger than T_1. We apply a technique of Breeden, Gibbons, and Litzenberger (1989) to construct mimicking portfolios of the factors so that we can perform the tests for the longer sample period.

Let f_c be the $T_1 \times 1$ vector that contains the time-series observations for a consumption-based factor (any of the *mcg* or *vcg* factors). Let R_1 be the $T_1 \times N$ matrix that contains the cross section of asset returns for the sample period. The weights of the mimicking portfolio are obtained from a standard regression of the factor on the asset returns:

$$w_{\text{mmk}} = \frac{w_c}{w_c' 1_N}, \quad w_c = (R_1' R_1)^{-1} R_1' f_c, \tag{16}$$

where 1_N is a $N \times 1$ vector of ones. Then, for the longer sample period, we obtain the mimicking portfolio version of the factor as

$$f_{\text{mmk}} = R_2 w_{\text{mmk}}, \tag{17}$$

where R_2 is the $T_2 \times N$ matrix of asset returns for the longer sample.

Tables VI and VII present estimation and test results for a sample constructed using this method that starts in July 1963 and ends in December 1995. The mimicking portfolios are constructed using the 1984–1995 CEX consumption data. Table VI is obtained using the 25 stock portfolios, and Table VII is obtained using the 28 assets, including the long-term bonds and the T-bill. The main findings are as follows. First, the results on the *vcg* factors largely confirm the results from the 1984–1995 period (see Tables IV and V). In Table VI the coefficients on the *vcg* factors are always significantly estimated with the expected sign. In Table VII this is only the case when the factors are based on cohorts (in rows 6 and 8). The HJ distances indicate significant improvements over the CAPM kernel estimated in row 13. Second, the estimation results on the *vcg* factors are robust to the inclusion of the Fama–French factors. Interestingly, it can be seen that the Fama–French factors are now estimated with the expected sign and are statistically significant in almost all cases. Again, this difference with the results in Tables IV and V is perhaps not surprising when

Table VI

Test Results Based on Mimicking Portfolios

This table presents results based on mimicking portfolios of consumption-based factors for the 1963–1995 period. It is otherwise identical to Table IV.

Row	Constant	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	J-Test	HJ-Distance
1	3.35	−2.28							142.74	0.725
	(4.88)	(−3.50)							(0.000)	
2	2.97			−1.92					143.74	0.728
	(4.96)			(−3.36)					(0.000)	
3	3.77	−1.80				−0.76	−5.25	−10.78	100.64	0.614
	(1.90)	(−2.54)				(−0.38)	(−3.16)	(−7.39)	(0.000)	
4	3.31			−1.27		−0.86	−5.10	−10.67	103.23	0.621
	(1.69)			(−2.05)		(−0.43)	(−3.10)	(−7.39)	(0.000)	
5	3.50	−3.01	0.62						134.25	0.676
	(4.88)	(−4.16)	(4.02)						(0.000)	
6	2.66			−1.91	0.32				142.12	0.707
	(4.40)			(−3.33)	(2.70)				(0.000)	
7	2.55	−3.09	0.66			0.99			133.12	0.675
	(2.13)	(−4.24)	(4.17)			(0.97)			(0.000)	
8	2.80			−1.89	0.32	−0.15			142.12	0.706
	(2.52)			(−3.24)	(2.70)	(−0.15)			(0.000)	
9	3.30	−2.64	0.56				−5.28	−9.74	93.96	0.576
	(4.39)	(−3.45)	(3.22)				(−5.00)	(−7.65)	(0.000)	
10	2.28			−1.52	0.43		−5.46	−10.49	97.86	0.591
	(3.68)			(−2.59)	(3.30)		(−5.33)	(−8.28)	(0.000)	
11	2.32	−2.81	0.59			1.10	−5.97	−9.33	94.45	0.574
	(1.11)	(−3.48)	(3.20)			(0.52)	(−3.42)	(−6.21)	(0.000)	
12	2.78			−1.42	0.43	−0.58	−5.12	−10.67	97.89	0.590
	(1.42)			(−2.27)	(3.23)	(−0.29)	(−3.09)	(−7.43)	(0.000)	
13	1.24					−0.26			157.38	0.753
	(1.21)					(−0.25)			(0.000)	
14	1.12						−5.55	−10.25	110.43	0.654
	(33.20)						(−5.34)	(−8.44)	(0.000)	
15	3.37					−2.19	−4.42	−11.09	106.48	0.650
	(1.73)					(−1.15)	(−2.82)	(−7.74)	(0.000)	

inspecting the raw asset returns in Table I. The book-to-market effect and size effect are clearly present in the 1963–1995 sample used in Table VI, but not in the 1984–1995 sample used in Table IV. The significance of the *vcg* factor seems to be robust to the inclusion of the size and book-to-market factors in both the 1984–1995 and the 1963–1995 samples. In summary, the results from the 1963–1995 sample period seem to confirm the importance of the consumption-based factors.

B.3. Exploring the Results

This section provides some more intuition and background behind the estimation and test results in Tables IV–VII, which verify the significance of the

Table VII

Test Results Based on Mimicking Portfolios with a Larger Set of Assets

This table presents results using the Fama–French 25 portfolios plus the two long-term bonds and the T-bill. It is otherwise identical to Table VI.

Row	Constant	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	J-Test	HJ-Distance
1	−2.68	3.70							178.10	0.815
	(−2.96)	(4.03)							(0.000)	
2	−3.41			4.46					176.50	0.802
	(−4.18)			(5.33)					(0.000)	
3	6.45	1.26				−6.41	−1.03	−12.39	109.00	0.678
	(3.57)	(0.83)				(−5.51)	(−0.60)	(−8.14)	(0.000)	
4	4.66			3.21		−6.54	0.39	−11.77	106.00	0.677
	(2.70)			(2.44)		(−5.60)	(0.24)	(−7.82)	(0.000)	
5	−2.56	3.70	−0.11						175.65	0.812
	(−2.84)	(4.05)	(−1.38)						(0.000)	
6	−4.54			4.94	0.73				159.30	0.763
	(−5.68)			(6.18)	(4.87)				(0.000)	
7	2.11	3.21	−0.09			−4.13			157.86	0.791
	(1.47)	(3.37)	(−1.16)			(−4.49)			(0.000)	
8	−0.64			4.18	0.62	−3.01			151.75	0.747
	(−0.45)			(4.91)	(4.17)	(−3.30)			(0.000)	
9	1.70	−0.56	0.00				−6.16	−10.55	134.86	0.734
	(1.18)	(−0.38)	(−0.04)				(−3.93)	(−7.52)	(0.000)	
10	−1.70			2.25	0.69		−3.65	−10.10	115.23	0.692
	(−1.35)			(1.81)	(4.56)		(−2.43)	(−7.25)	(0.000)	
11	6.59	1.06	0.04			−6.41	−1.20	−12.40	109.12	0.678
	(3.56)	(0.68)	(0.42)			(−5.49)	(−0.69)	(−8.19)	(0.000)	
12	3.28			3.31	0.56	−5.79	0.80	−11.18	99.45	0.636
	(1.89)			(2.55)	(3.60)	(−4.93)	(0.50)	(−7.54)	(0.000)	
13	5.36					−4.29			161.23	0.803
	(5.99)					(−4.93)			(0.000)	
14	1.13						−5.79	−10.34	135.61	0.736
	(33.44)						(−5.49)	(−8.42)	(0.000)	
15	7.58					−6.29	−1.94	−12.82	108.53	0.678
	(6.36)					(−5.46)	(−1.58)	(−9.19)	(0.000)	

consumption-based factors in the presence of other factors in the pricing kernel. Table VIII provides a more formal approach to the question of whether one (set of) factor(s) drives out another. We follow Cochrane (1996) in reporting test statistics for excluding certain factors from a pricing kernel that nests several existing models.[19] In rows 1–6, the unrestricted model consists of five factors: The two consumption-based factors constructed using cohorts and the three

[19] Several recent papers (e.g., see Lettau and Ludvigson (2002)) insert variables like size and book-to-market in an asset pricing relationship as a specification test, to check if there are residual effects unexplained by the factors. The purpose of our regressions is slightly different, because we do not use firm-specific information but rather factors. The objective of the exercise is rather to verify whether different factors exhibit collinearity, much as in Fama and French (1993).

Table VIII

Model Comparison Tests

This table presents results from model comparison tests. The tests are based on ΔJ_T, which is the difference between the J_T statistic of the restricted model (that excludes a given set of factors) and that of the unrestricted model. See Cochrane (1996) for further details. Panel A is for the 1984–1995 sample period, using two sets of test assets: The Fama–French 25 portfolios as the first set (FF25) and the 25 portfolios plus the two long-term bonds and T-bill as the second set (FF25 + 3). Panel B is for the 1963–1995 period with the two sets of assets. In each case, "×" denotes an excluded factor for the restricted model. For row 1 through 6, the unrestricted model consists of five factors: Two consumption-based factors (mcg_{ae} and vcg_{ae}) and the three Fama–French factors (R_M, SMB, and HML). For rows 7 and 8, the unrestricted model consists of the two consumption-based factors. The ΔJ_T statistics and the corresponding p-values are reported.

	Factors					FF25		FF25 + 3	
Row	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	ΔJ_T	p-Value	ΔJ_T	p-Value
					Panel A: 1984–1995				
1	×	×				35.29	0.0000	24.92	0.0000
2	×					20.80	0.0000	7.81	0.0052
3		×				22.64	0.0000	24.91	0.0000
4			×			7.82	0.0052	3.21	0.0730
5				×	×	14.09	0.0009	20.68	0.0000
6			×	×	×	22.03	0.0001	21.52	0.0001
	mcg_{ae}	vcg_{ae}							
7	×					18.39	0.0000	1.16	0.2822
8		×				27.08	0.0000	27.83	0.0000
					Panel B: 1963–1995				
1	×	×				14.04	0.0009	18.69	0.0001
2	×					5.16	0.0232	6.52	0.0107
3		×				10.45	0.0012	12.95	0.0003
4			×			0.08	0.7741	24.26	0.0000
5				×	×	80.36	0.0000	65.59	0.0000
6			×	×	×	80.64	0.0000	68.46	0.0000
	mcg_{ae}	vcg_{ae}							
7	×					11.12	0.0009	38.23	0.0000
8		×				7.30	0.0069	23.72	0.0000

Fama–French factors. For the 1984–1995 sample in Panel A, the test statistics indicate that the data do not support excluding the consumption-based factors from the regression. On the other hand, the data also do not support excluding the Fama–French factors even in the presence of the consumption-based factors. Panel B confirms that this is also the case for the 1963–1995 period, even though the statistical evidence is less strong than in the 1984–1995 period. To put this in perspective, it is useful to note that dropping the market factor from the five-factor pricing kernel is not rejected at conventional significance levels for the 1963–1995 sample.

While the iterated GMM-based estimation results are asymptotically efficient and the HJ distances based on one-stage estimation can be readily compared across pricing kernels, some researchers prefer comparing models using

Table IX

Descriptive Statistics and Goodness-of-Fit Measures

This table presents descriptive statistics for several factor models. Panel A and B are for the 1984–1995 sample period, using two sets of test assets: The Fama–French 25 portfolios as the first set (FF25) and the 25 portfolios plus the two long-term bonds and T-bill as the second set (FF25 + 3). Panel C and D are for the 1963–1995 period with the two sets of assets. For each model, we report the mean and standard deviation of the stochastic discount factor (SDF), the HJ distance, and the adjusted R-square $\bar{R}^2$ of the cross-sectional regression (OLS and GLS). The SDFs are evaluated using the parameter estimates that minimize the HJ-distance measure. The consumption factors are constructed with the age-and-education cohorts derived from the nondurable and services consumption for asset holders. That is, the consumption factors are the pair mcg_{ae} and vcg_{ae} described in Appendix B.

SDF			$\bar{R}^2$		SDF			$\bar{R}^2$	
Mean	SD	HJ-Distance	OLS	GLS	Mean	SD	HJ-Distance	OLS	GLS
				Panel A: 1984–1995, FF25					
		Market Factor (R_m)					Fama–French Three Factors		
0.938	0.625	1.892	0.371	0.458	0.937	0.685	1.880	0.484	0.413
		mcg_{ae}					mcg_{ae} and vcg_{ae}		
0.941	0.993	1.889	0.009	0.481	0.942	1.627	1.791	0.320	0.530
				Panel B: 1984–1995, FF25 + 3					
		Market Factor (R_m)					Fama–French Three Factors		
0.994	0.317	2.743	−0.038	0.173	0.994	0.592	2.699	0.518	0.133
		mcg_{ae}					mcg_{ae} and vcg_{ae}		
0.994	0.186	2.759	0.004	0.188	0.995	1.427	2.696	0.079	0.220
				Panel C: 1963–1995, FF25					
		Market Factor (R_m)					Fama–French Three Factors		
0.984	0.013	0.753	−0.038	−0.051	0.989	0.406	0.650	0.807	0.153
		mcg_{ae}					mcg_{ae} and vcg_{ae}		
0.988	0.219	0.728	0.184	0.014	0.987	0.276	0.707	0.367	0.052
				Panel D: 1963–1995, FF25 + 3					
		Market Factor (R_m)					Fama–French Three Factors		
0.998	0.222	0.803	0.266	0.118	0.998	0.488	0.678	0.892	0.318
		mcg_{ae}					mcg_{ae} and vcg_{ae}		
0.998	0.223	0.802	0.426	0.137	0.998	0.333	0.763	0.442	0.145

different criteria. Table IX reports adjusted R^2s obtained using cross-sectional regressions for the CAPM, the Fama–French three-factor model and the one-factor and two-factor consumption-based models.[20] Panels A and B report results for the 1984–1995 sample, and Panels C and D for the 1963–1995 sample. Panel A reports on the case with 25 stock returns, and Panel B augments this data set with the long-term bonds and the T-bill. We present adjusted R^2s using

[20] Kan and Zhang (1999) document some problems associated with two-pass tests of asset pricing models such as the cross-sectional regressions reported here. However, the R^2 from such tests remains a valid and widely used measure for goodness of fit.

OLS as well as GLS estimates, and there are some differences between the two sets of results, which confirms existing results in the literature (see Cochrane (2001)). The OLS-based results are particularly simple to interpret because they are based on equally weighted pricing errors.

We are interested in the relative rankings between the models, and these results are relatively robust. The adjusted R^2s for the CAPM are relatively low, similar to the results obtained by Jagannathan and Wang (1996) and Lettau and Ludvigson (2002). While the one-factor consumption-based models also yield low adjusted R^2s, the two-factor consumption-based models provide a better fit. It must be noted that the overall low values of the adjusted R^2s are not relevant. They simply reflect that the cross section under study here is extremely challenging, which also results in all models having high HJ distances. The only relevant interpretation of the adjusted R^2s is as an indicator of the differences in fit between models. Finally, it must be noted that whereas the Fama–French three-factor model performs worse than the consumption-based models when judged by the HJ distance, this is not the case when judged by means of the adjusted R^2. These findings are not necessarily surprising, because the weighting of the assets used in the computation of HJ distances and adjusted R^2s in cross-sectional regressions is very different. In Table IX, we attempt to surmise why the consumption-based models do better than the three-factor Fama–French model when using the HJ distance. The HJ distance associated with a constant pricing kernel is 2.762 (not reported). This means that a constant pricing kernel is a very bad candidate in this case and that volatility in the pricing kernel will to some extent be rewarded (if those movements are correlated with changes in returns). Table IX indicates that while the standard deviation of the Fama–French kernel is not much higher than the standard deviation of the CAPM pricing kernel, the kernels associated with the two-factor consumption-based models are more variable. Panels C and D of Table IX confirm the results of Panels A and B, even though in some instances the CAPM and the one-factor consumption-based models perform better in this sample. Also, the standard deviations of the different pricing kernels do not differ as dramatically as in Panels A and B.

Figures 1 and 3 graphically illustrate the OLS-based results in Table IX. The pictures on the left represent the case of 25 assets. The pictures on the right represent the case of 28 assets. For the 25 Fama–French portfolios, each two-digit number in these figures represents one portfolio. The first digit refers to size, with 1 indicating the portfolio with the smallest-size firms. The second digit refers to book-to-market, with 1 indicating low book-to-market. The T-bill, corporate bond, and government bond are indexed by *tb*, *cb*, and *gb*, respectively. The figures clearly indicate that the performance of the two-factor consumption-based model is far superior to that of the CAPM, but that with equally weighted pricing errors, the Fama–French three-factor model outperforms the consumption-based model. These figures also allow a number of other observations that help us to understand the performance of the consumption-based model. First, the same portfolios pose the most stringent challenges to all models under consideration. This is most clearly seen in the 1984–1995

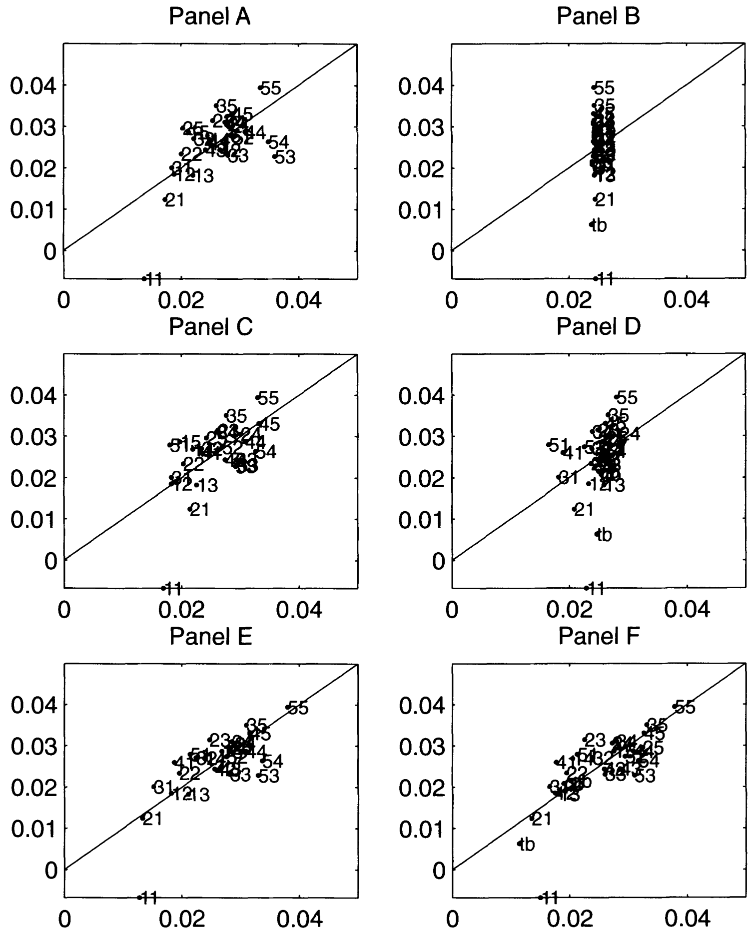

Figure 1. Cross-sectional regressions (1984–1995). Mean returns (vertical axes) are plotted against predicted mean returns (horizontal axes) based on OLS cross-sectional regressions for the market factor (Panels A and B), the *mcg* and *vcg* factors (Panels C and D), and the Fama–French three factors (Panels E and F), for the sample period 1984–1995. Panels A, C, and E are based on OLS estimation using the Fama–French 25 portfolios. Panels B, D, and F are based on the Fama–French 25 portfolios plus the two long-term bonds and the T-bill. For the 25 Fama–French portfolios, each two-digit number represents one portfolio. The first digit refers to size, with 1 indicating the portfolios with the smallest-size firms. The second digit refers to book-to-market, with 1 indicating low book-to-market. The T-bill, corporate bond, and government bond are indexed by *tb*, *cb*, and *gb*, respectively.

period, where the small size/low book-to-market portfolio (indexed 11) is by far the largest outlier for all three models, but it can be seen from the figures that this is also the case for other outliers. Second, it can be seen especially for the 1963–1995 sample that the consumption-based model has some problems in explaining the return on the T-bill, and to some extent also the return on the bond portfolios. This is perhaps surprising, because intuitively one might

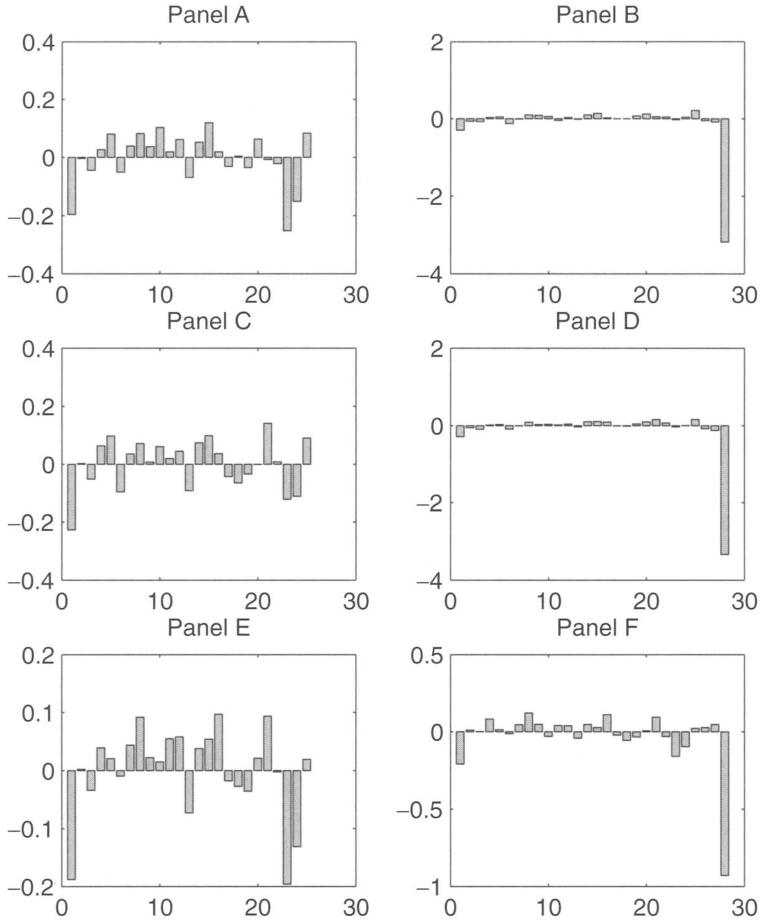

Figure 2. Scaled pricing errors (1984–1995). Plotted in each panel are the pricing errors from the corresponding panel of Figure 1, scaled by the asset return standard deviation. For Panels A, C, and E, the Fama–French 25 portfolios are indexed by the numbers 1–25 on the horizontal axes. For Panels B, D, and F, the two long-term bonds and the T-bill are added; they are indexed by 26, 27, and 28.

expect the consumption-based model to perform relatively well in this dimension. Third, in both samples the signs of the pricing errors for the different assets are usually the same across the three models. This indicates that the consumption-based and Fama–French models address the shortcomings of the CAPM, without completely eliminating them or overturning the pricing biases present in that model. Interestingly, Lettau and Ludvigson (2002) arrive at some of the same conclusions for their model.

Finally, Figures 2 and 4 provide some more background on the properties of the T-bill return, to address the problems of the consumption-based model to explain this return. Each panel provides pricing errors scaled by the

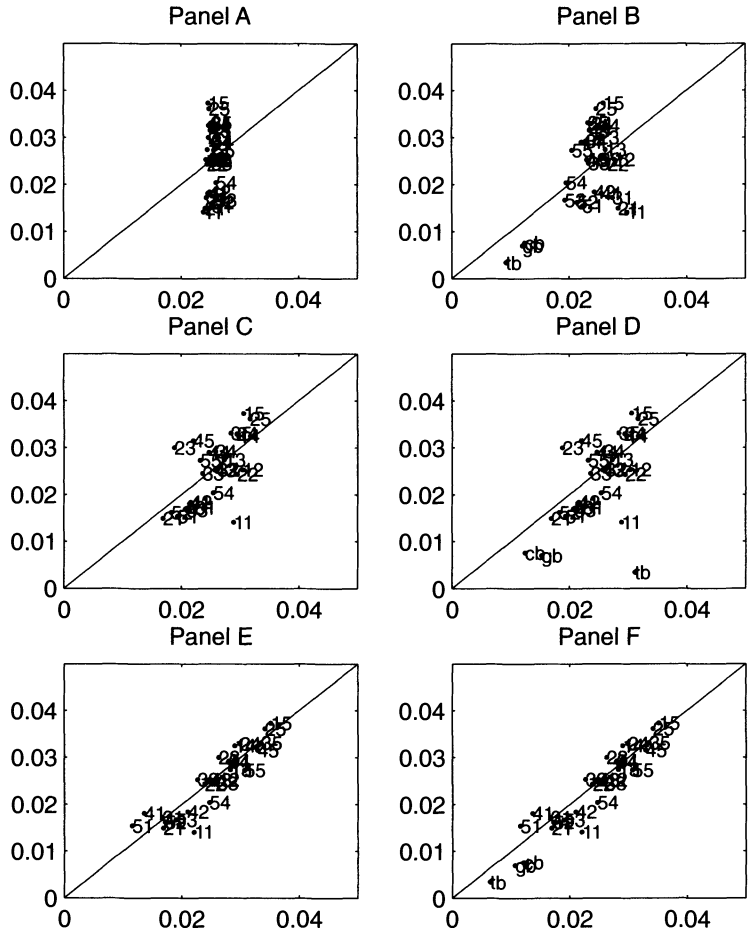

Figure 3. Cross-sectional regressions (1963–1995). Mean returns (vertical axes) are plotted against predicted mean returns (horizontal axes) based on OLS cross-sectional regressions for the market factor (Panels A and B), the *mcg* and *vcg* factors (Panels C and D), and the Fama–French three factors (Panels E and F), for the sample period 1963–1995. Panels A, C, and E are based on OLS estimation using the Fama–French 25 portfolios. Panels B, D, and F are based on the Fama–French 25 portfolios plus the two long-term bonds and the T-bill. For the 25 Fama–French portfolios, each two-digit number represents one portfolio. The first digit refers to size, with 1 indicating the portfolios with the smallest-size firms. The second digit refers to book-to-market, with 1 indicating low book-to market. The T-bill, corporate bond, and government bond are indexed by *tb*, *cb*, and *gb*, respectively.

corresponding standard deviation. Note that when minimizing the HJ distance, this is the metric that is being minimized. The panels on the left contain the pricing errors for the case of 25 assets. It can be seen that even though there are some outliers, they are not extreme. The panels on the right contain the pricing errors for the case of 28 assets. It is clear that in all cases, there is one extreme outlier, which corresponds to the T-bill. While this does not explain

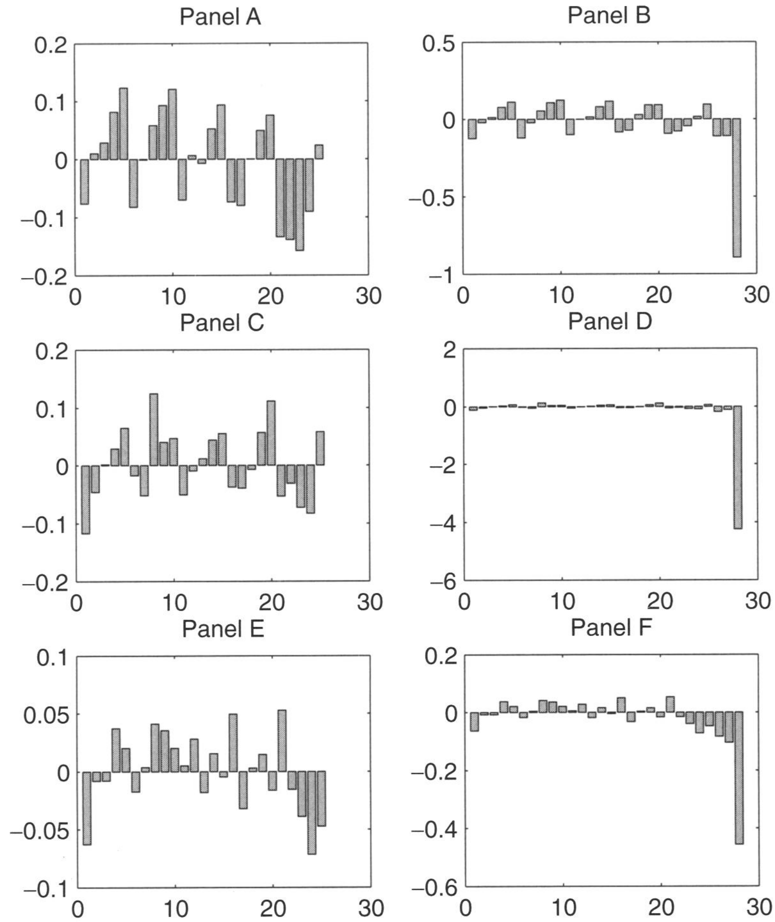

Figure 4. Scaled pricing errors (1963–1995). Plotted in each panel are the pricing errors from the corresponding panel of Figure 3, scaled by the asset return standard deviation. For Panels A, C, and E, the Fama–French 25 portfolios are indexed by the numbers 1–25 on the horizontal axes. For Panels B, D, and F, the two long-term bonds and the T-bill are added; they are indexed by 26, 27, and 28.

why the *mcg* factor sometimes shows up with the wrong sign in Table V, it does suggest that the properties of the T-bill are quite different from those of the other assets, and therefore it may not be surprising that its inclusion in the cross section significantly affects estimation results.

B.4. Test Results Using the CPS Income Factors

Table X presents estimation results obtained using the CPS income growth factors. The results are mixed. The pricing kernel with the income growth factors in row 4 yields a significantly estimated positive sign on the variance factor *vig*, but the estimate on *mig* has the wrong sign. When combining the

Table X

Testing for the Significance of CPS Income Factors

This table includes results for seven linear pricing kernels, which are various combinations of the two CPS income factors and the Fama–French three factors. The income factors, mig_t and vig_t, are the cross-sectional mean and variance of income growth constructed using 34 annual observations from the Consumer Population Survey (1967–2000). The variable R_M is the market return, and SMB and HML are the size and book-to-market factors of Fama and French. Our test assets consist of the 25 Fama–French size and book-to-market portfolios. A standard GMM procedure is implemented for testing the moment conditions $E[M_t(\beta)R_{it}] = 1$, where R_{it} is the return on the i^{th} test asset. In the initial round, the HJ-distance measure is minimized. Then the iterated GMM estimates are obtained; that is, at each round, the weighting matrix is updated using the estimates from the previous round, and the procedure is repeated until estimates converge. Reported in the table are the iterated estimates and the J-test statistics that are based on the iterated estimates. In parentheses under the estimates are t-statistics and in parentheses under the J statistics are the p-values. The HJ distances are also included.

Row	Const.	*mig*	*vig*	R_M	SMB	HML	J-Test	HJ-Distance
1	0.95			−0.30			26.70	1.77
	(10.08)			(−0.24)			(0.319)	
2	0.97				−0.78	−0.77	27.46	1.75
	(11.54)				(−0.65)	(−0.69)	(0.237)	
3	1.18			−2.01	−0.97	−1.87	28.65	1.74
	(6.93)			(−1.62)	(−0.85)	(−1.49)	(0.155)	
4	−12.97	11.97	6376.99				30.38	1.76
	(−2.00)	(1.88)	(10.43)				(0.139)	
5	4.32	−3.57	661.48	−1.80			29.34	1.76
	(0.53)	(−0.45)	(1.44)	(−1.63)			(0.135)	
6	6.99	−6.06	265.16		−2.06	−0.56	28.29	1.75
	(0.72)	(−0.62)	(0.64)		(−1.43)	(−0.52)	(0.132)	
7	5.83	−4.39	−131.09	−3.24	−1.41	−2.16	29.05	1.74
	(0.77)	(−0.58)	(−0.47)	(−2.39)	(−1.01)	(−1.69)	(0.087)	

income-based factors with other factors in rows 5–7, the variance factor is not significantly estimated. The HJ distances for these pricing kernels do not differ very much from the HJ distance for the CAPM. We conclude that the results using income-based CPS factors are unconvincing. It is likely that the small size of the sample (34 annual observations) is responsible for this.

IV. Concluding Remarks

A number of recent papers have demonstrated that the presence of uninsurable idiosyncratic consumption risk is relevant to the explanation of well-established puzzles in the asset-pricing literature, such as the equity premium puzzle and the risk-free rate puzzle. This paper shows that the presence of such risk is also useful for constructing factors that help explain the cross section of asset returns. We investigate the performance of a pricing kernel linear in the first and the second moments of the cross-sectional distribution of consumption growth. We find that it is imperative to address the presence of measurement error in consumption by constructing synthetic cohorts. Using the consumption factors based on synthetic cohorts, we find that the consumption-based pricing factors are almost always significantly estimated. Parameter estimates on the first moment of the cross-sectional distribution usually indicate that assets are riskier if positively correlated with consumption growth, and parameter estimates on the consumption dispersion factor most often indicate that assets are riskier if negatively correlated with consumption dispersion. This negative correlation is also of interest because it is known to be helpful in resolving the equity premium puzzle (see Constantinides and Duffie (1996)).

We evaluate the importance of this finding by comparing the pricing performance of the consumption-based factors with some well-established benchmarks, using HJ distances and R^2s as yardsticks. Our original sample period is fairly short because of the limited availability of consumption data, but we also investigate the robustness of our results by extending the sample period using mimicking portfolios. We find that the two-factor consumption-based kernels outperform the CAPM over the sample period. Also, the performance of the two-factor consumption-based kernels compares favorably with that of the Fama–French three-factor model. When estimating kernels that combine the CAPM or Fama–French factors with the consumption-based factors, the consumption-based factors are still estimated significantly with the same signs. Statistical tests do not advocate dropping these factors from kernels that nest these different pricing models. Alternatively, these tests also do not advocate dropping the Fama–French factors from these kernels.

The traditional view is that consumption-based models are not very helpful for cross-sectional asset pricing. However, in a recent paper, Lettau and Ludvigson (2002) demonstrate that conditional versions of consumption-based models perform much better than unconditional versions. This paper provides further evidence that the empirical performance of consumption-based models is probably more satisfactory than is commonly thought. Given that these factors are suggested by theory within the context of a well-specified multiperiod model, they deserve to be given close attention. To build an even stronger case, two exercises related to the ones in this paper come to mind. First, the analysis in this paper is limited to unconditional models. Given the analysis in Lettau and Ludvigson (2002), an extension to conditional models may be of interest. Second, an extension of the analysis in this paper to pricing models that incorporate higher moments of the distribution of consumption growth may prove interesting.[21]

Appendix A. A Linear Approximation to the Pricing Kernel

Using the methodology of Mankiw (1986), Cogley (2002) and Brav et al. (2002), we approximate $M(cg_{it})$ with a second-order polynomial

$$\frac{1}{N}\sum_{i=1}^{N} M(cg_{it}) \approx M(mcg_t) + \frac{1}{2}M''(mcg_t)vcg_t, \tag{A1}$$

[21] This analysis may be difficult, especially in the case of the cohort data. The cross section contains relatively limited information to estimate higher moments accurately.

where mcg_t is the mean of consumption growth and vcg_t is the variance of consumption growth. Here we consider the TS-CRRA with per period utility function $c_{it}^{1-\alpha}/(1-\alpha)$, but the analysis can be extended to other utility functions. For the TS-CRRA utility function, the pricing kernel is

$$M(cg_{it}) = \delta cg_{it}^{-\alpha}, \tag{A2}$$

where δ is the time discount rate. The approximation then becomes

$$\frac{1}{N}\sum_{i=1}^{N} M(cg_{it}) \approx \delta mcg_t^{-\alpha} + \frac{1}{2}\delta\alpha(\alpha+1)mcg_t^{-\alpha-2}vcg_t. \tag{A3}$$

Now consider first-order Taylor series expansions for both terms involving mcg_t around $mcg_t = 1$. We get

$$\delta mcg_t^{-\alpha} \approx \delta - \alpha\delta(mcg_t - 1) \tag{A4}$$

$$\frac{1}{2}\delta\alpha(\alpha+1)mcg_t^{-\alpha-2} \approx \frac{1}{2}\delta\alpha(\alpha+1) - \frac{1}{2}\delta\alpha(\alpha+1)(\alpha+2)(mcg_t - 1) \tag{A5}$$

and therefore

$$\begin{aligned}\frac{1}{N}\sum_{i=1}^{N} M(cg_{it}) &\approx \delta - \alpha\delta(mcg_t - 1) + \frac{1}{2}\delta\alpha(\alpha+1)vcg_t \\ &- \frac{1}{2}\delta\alpha(\alpha+1)(\alpha+2)(mcg_t - 1)vcg_t.\end{aligned} \tag{A6}$$

Ignoring the term in $(mcg_t - 1)vcg_t$, we get our proposed pricing kernel

$$\frac{1}{N}\sum_{i=1}^{N} M(cg_{it}) = M_t \approx \beta_0 + \beta_1 mcg_t + \beta_2 vcg_t, \tag{A7}$$

where the estimates β_0, β_1, and β_2 can be interpreted as combinations of the behavioral parameters δ and α. Whether it is realistic to ignore the term in $(mcg_t - 1)vcg_t$ is open to debate. Given that $(mcg_t - 1)$ and vcg_t are small, we can ignore it if the rate of risk aversion is not too large. The representative agent literature argues that the rate of risk aversion is large (see Mehra and Prescott (1985)). Part of the incomplete markets literature finds a much smaller rate of risk aversion (see Jacobs (1999) and Brav, Constantinides, and Geczy (2002)).

Appendix B. Data Construction

B.1. CEX Consumption Data

The CEX is not a genuine panel data set, but a series of cross sections with a limited time dimension. However, in the context of the exercise proposed in this paper, this is not necessarily a very serious problem, because at each time we simply use every available cross section to construct cross-sectional moments. The advantage of the CEX is that it provides very extensive data on consumption. Our empirical tests use a measure of NDS consumption constructed from the CEX. NDS consumption is constructed as the sum of expenditures on food, alcoholic beverages, tobacco, gas, utilities, apparel, public transportation, household operations, and personal care. Table II indicates that NDS consumption is low compared to NIPA data (not reported). The reason is that the construction of services from CEX data on household consumption is not straightforward. As a result, the NDS data in this paper and in other studies that use the CEX do not contain important consumption categories. A related observation is that in many areas the distinction between durable and nondurable consumption is tenuous. For instance, clothing is included in (the narrow definition of) NDS consumption in this paper. See also Hayashi (1985) and Mankiw, Rotemberg, and Sunmers (1985) on this matter.

We conduct our analysis for a sample that only contains asset holders, and also for a sample that contains all households. While the sample including asset holders is well motivated by theory, the sample including all households is not. We use that sample as a robustness exercise, because it is much larger and therefore more reliable when constructing cohorts. To classify an individual as an asset holder, the literature has often identified stock holders as asset holders. Unfortunately, the CEX does not directly ask if an individual holds stocks either directly or indirectly through a pension plan. We therefore proceed to construct a sample of households that are likely to hold stocks. It consists of households that report the existence of at least one of the following: (i) holdings of stocks or bonds, (ii) dividend income, and/or (iii) contributions to an IRA. In the CEX, these questions are asked in reference to two points in time, the first and the last (fifth) quarter that the households are in the sample. To determine which households are asset holders, we use the answer referring to the first quarter. Using this classification, the percentage of asset holders is approximately 31%, which is comparable to the number in Mankiw and Zeldes (1991) but lower than the number in Cogley (2002).

The structure of the CEX is special and deserves comment. The CEX data are organized in quarterly files and participants' interviews are scheduled 3 months apart, but these interviews are staggered evenly throughout the year on a monthly basis. That is, if the total number of participants in the survey is X in a given year, then every month $X/12$ participants are interviewed and asked about their consumption in the previous 3 months. As the same household is interviewed several times, this allows us to construct cross sections of consumption growth rates based on 3-month consumption, but at a monthly frequency. Vissing-Jorgensen (2002) and Cogley (2002) also use the CEX to construct consumption growth at a monthly frequency. Brav, Constantinides, and Geczy (2002) use a slightly different approach. Using this construction, the first available time period to construct consumption growth rates using the CEX data is April 1984. Due to a data matching problem, we cannot use data for the first 4 months of 1986. Moreover, several indicators revealed low data quality for the last 3 months available (October through December 1995). We therefore exclude these months, which leaves us with 134 monthly cross sections on consumption growth.

There are a few complications due to the limitations of the data. The CEX reports consumption for the household unit. We attempt to correct for household size using two techniques. First, we simply divide household consumption by the number of household members. Alternatively, we use a more sophisticated approach by correcting for family size using a scale that is estimated from the data. As both approaches do not significantly affect the test results, we present results using household consumption as the unit of observation. Another complication is the presence of seasonalities. It is well known that seasonalities are present in consumption data and that they are important for asset pricing (see Miron (1986) and Ferson and Harvey (1992)). When inspecting the raw CEX household consumption data, seasonalities seem to be even more pronounced than in NIPA data. The most obvious manifestation of this finding is the well-known decrease in consumption in the first quarter. Reported results do not adjust for seasonality, in accordance with other papers that use the CEX (see Attanasio and Weber (1995), Brav, Constantinides, and Geczy (2002), Vissing-Jorgensen (2002), Cogley (2002), and Balduzzi and Yao (2000)). We performed a robustness exercise by controlling for seasonality using the census X11 method as implemented in EVIEWS. Even though the resulting seasonal adjustment factors are significant (as is the case for NIPA data), they do not affect test results.

B.2. Asset Return Data

We use monthly return data from Fama and French (1993) to construct time series of returns that are consistent with the structure of the 3-month consumption growth data described in Appendix B.1. To illustrate, consider how this is done for a particular month, say April 1990. In an April 1990 interview, the consumption refers to consumption in January, February, and March 1990 and lagged consumption refers to consumption in October, November, and December 1989. To match returns with these consumption data, we first compute three 3-month rolling returns for October 1989 to January 1990, November 1989 to February 1990, and December 1989 to March 1990. We then use the arithmetic average of these three returns. Note that even though the consumption and consumption growth data are for 3-month periods, the frequency of the consumption cross sections is monthly; therefore, this is also the case for the returns data. As a result, the returns data we use in our empirical tests are partly overlapping.

The Fama–French portfolios are now widely used. They are value-weighted portfolios of stocks listed on the NYSE, AMEX, and Nasdaq. These portfolios are sorted on firm size and book-to-market equity and exhibit strong cross-sectional dispersion in average returns. For more details on the portfolios, see Fama and French (1993). For the bond returns, we use the total return on Treasury bonds (the CRSP variable GBTRET), the total return on long-term corporate bonds (the CRSP variable CBTRET), and the 3-month T-bill rate. For the market portfolio, we use the CRSP value-weighted portfolio of stocks listed on NYSE, AMEX, and Nasdaq that Fama and French use to proxy for the market portfolio. We follow a long tradition in the finance literature by measuring $R_{M,t}$ using the return on publicly listed stocks. A large number of studies have debated whether to include the return on human capital in this construction (see Mayers (1972), Roll (1977), Fama and Schwert (1977), Campbell (1996), Jagannathan and Wang (1996), and Lettau and Ludvigson (2002)). We do not analyze this issue because our primary motivation for studying the CAPM here is to find an appropriate benchmark for the performance of our consumption-based model, and not the appropriate specification of the market return. Also included in our empirical tests are the size (SMB) and book-to-market (HML) factors of Fama and French. All the variables are in real terms.

B.3. Synthetic Cohort Construction

The synthetic cohorts approach involves the construction of a representative consumer for a typical group that can be defined by observable characteristics such as age. It is clear that for most plausible parameterizations of measurement error this construction will mitigate its effects, without eliminating them. It must also be noted that the motivation for this technique is of course very similar to the motivation for testing the CAPM using portfolios instead of individual assets, as originally implemented by Black, Jensen, and Scholes (1972) and Fama and MacBeth (1973).

Unfortunately, the choice of grouping method for the construction of synthetic cohorts is not obvious. On the one hand, one does not want the groups to be too small, because in that case the effects of measurement error are not likely to disappear. On the other hand, by making the groups too large, it is clear that one constructs away the potential impact of idiosyncratic consumption risk. It is not obvious that there is a realistic optimal solution to this problem. The optimal choice depends on the size and the type of the measurement error, and by definition we do not know a lot about this. We construct synthetic cohorts based on two very simple grouping variables: the age and the education of the head of the household. To understand the problems implied by this choice, note that for all specifications we work with two samples, one with all households and another with asset holders only. It is clear that the choice to hold assets depends critically on age and education. Therefore, the composition and size of a given cohort will be different in both samples, and this could influence test results. To minimize (potential) problems with small cohorts, we impose a constraint on the cohort construction: We include only individuals older than 25 and younger than 60 in the sample. We then proceed to construct factors in the following two ways. One set of factors is based on age only: The head of the household in our sample can be any of 36 different ages (between 25 and 60). We use this to split the sample into 12 age brackets, or cohorts, and subsequently construct the pricing factors in each quarter using 12 observations (cohorts). Presuming that the consumption patterns of households of similar ages are comparable, cohorts are based on age proximity; for example, the first cohort consists of household heads aged 25, 26, and 27. For most of our empirical results, we use a set of factors based on age as well as education as a sorting variable. We use the educational information in the CEX to create a sample of consumers who have at least completed a college education, and another sample of consumers who have not. This gives us two educational categories, which we combine with six age categories to again get 12 cohorts.

We refer to the construction of factors using individual data with the i subscript, that is

$$mcg_{i,t} = \frac{1}{N}\sum_{i=1}^{N}(cg_{i,t}), \quad vcg_{i,t} = \frac{1}{N}\sum_{i=1}^{N}(cg_{i,t} - mcg_{i,t})^2, \tag{B1}$$

where $cg_{i,t} = (c_{i,t}/c_{i,t-1})$ and $c_{i,t}$ is the consumption of individual i at time t. To construct factors based on cohort consumption growth, we aggregate within a cohort on the level of consumption. We also experimented extensively with aggregating on consumption growth (instead of the consumption level). Test results obtained using this construction are not robust. This finding is consistent with the existence of certain types of measurement error in consumption (such as additive measurement error) that can be removed by aggregating on consumption but not necessarily by aggregating on consumption growth. The consumption of a representative cohort k at time t is given by $\mathrm{cohc}_{k,t} = (1/N_{k,t})\sum_{i=1}^{N_{k,t}}(c_{i,t})$, where $N_{k,t}$ is the number of observations on this cohort at time t. The cohort consumption growth can then be denoted by $\mathrm{cohcg}_{k,t} = (\mathrm{cohc}_{k,t}/\mathrm{cohc}_{k,t-1})$. For example, if the cohort is based on age, the factors can then be computed as

$$mcg_{a,t} = \frac{1}{H}\sum_{k=1}^{H}(cohcg_{a,k,t}), \quad vcg_{a,t} = \frac{1}{H}\sum_{k=1}^{H}(cohcg_{a,k,t} - mcg_{a,t})^2, \tag{B2}$$

where H is the number of cohorts, and the subscript a denotes the cohort construction that is based on age.

To address the problem that not all cohorts contain the same number of observations, we investigate the robustness of our results using an alternative construction. Remember that $N_{k,t}$ denotes the number of households in cohort k at time t and N_t the total number of households at time t. Define $w_{k,t} = N_{k,t}/N_t$. We then define the alternative factors as

$$\begin{aligned} mcg_{a,Alt,t} &= \sum_{k=1}^{H} w_{k,t}(cohcg_{k,t}), \\ vcg_{a,Alt,t} &= \sum_{k=1}^{H} w_{k,t}(cohcg_{k,t} - mcg_{a,Alt,t})^2, \end{aligned} \tag{B3}$$

where $\mathrm{cohcg}_{k,t} = (\mathrm{cohc}_{k,t}/\mathrm{cohc}_{k,t-1})$, $\mathrm{cohc}_{k,t} = (1/N_{k,t})\sum_{i=1}^{N_{k,t}}(c_{i,t})$, and H is the number of cohorts. In other words, these alternative factors use the same cohort information but weigh the results according to the number of households in each cohort. When we repeat the analysis with these alternative factors, our conclusions are not affected.

Appendix C. Some Robustness Results

This Appendix contains additional robustness checks. The results are reported in Tables CI–CIII.

Table CI

Test Results with Age Cohorts

This table presents results for the case where the consumption-based factors are constructed with age cohorts. It is otherwise identical to Table IV.

Row	Const.	mcg_i	vcg_i	mcg_a	vcg_a	R_M	SMB	HML	J-Test	HJ-Distance
1	28.80	−26.34							81.80	1.905
	(6.03)	(−5.88)							(0.000)	
2	23.64			−22.69					88.43	1.887
	(6.19)			(−6.03)					(0.000)	
3	24.71	−33.08				10.88	−7.64	−11.11	69.92	1.778
	(4.09)	(−5.54)				(2.49)	(−1.51)	(−2.54)	(0.000)	
4	18.96			−34.09		15.56	−12.40	−10.49	61.88	1.753
	(3.19)			(−5.55)		(3.36)	(−2.36)	(−2.46)	(0.000)	
5	30.64	−28.69	3.99						79.24	1.863
	(6.07)	(−5.84)	(1.01)						(0.000)	
6	15.50			−16.58	196.52				94.10	1.865
	(4.89)			(−5.03)	(4.03)				(0.000)	
7	24.47	−31.44	4.64			8.81			73.12	1.746
	(4.30)	(−5.69)	(1.07)			(2.65)			(0.000)	
8	2.00			−24.13	397.93	18.76			57.81	1.724
	(0.34)			(−4.30)	(4.37)	(5.22)			(0.000)	
9	31.61	−29.27	2.12				0.91	−9.23	77.32	1.819
	(6.03)	(−5.78)	(0.63)				(0.23)	(−2.28)	(0.000)	
10	28.39			−28.93	156.23		4.05	−14.49	79.92	1.798
	(6.23)			(−6.35)	(2.74)		(1.15)	(−3.68)	(0.000)	
11	26.10	−37.63	9.33			12.73	−7.11	−11.44	66.67	1.693
	(4.10)	(−5.37)	(1.33)			(2.50)	(−1.34)	(−2.59)	(0.000)	
12	2.49			−33.27	516.86	26.16	−20.76	−5.83	38.74	1.630
	(0.31)			(−4.20)	(4.01)	(4.38)	(−2.59)	(−1.03)	(0.005)	
13	−5.65					6.48			108.01	1.892
	(−3.86)					(4.48)			(0.000)	
14	1.00						2.34	−3.00	109.55	1.950
	(21.47)						(1.05)	(−1.68)	(0.000)	
15	−16.64					17.18	−20.20	−2.66	92.84	1.880
	(−5.67)					(5.99)	(−6.25)	(−1.22)	(0.000)	

Table CII

Test Results with All Households

This table presents results for the case where the consumption-based factors are constructed with consumption data of all households. It is otherwise identical to Table IV.

Row	Const.	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	J-Test	HJ-Distance
1	29.25	−26.65							87.15	1.876
	(6.47)	(−6.34)							(0.000)	
2	24.71			−23.88					92.94	1.868
	(6.39)			(−6.26)					(0.000)	
3	23.26	−40.64				19.87	−20.30	−13.82	70.24	1.748
	(3.28)	(−5.37)				(3.66)	(−3.10)	(−3.19)	(0.000)	
4	20.28			−38.65		18.36	−20.14	−11.46	59.02	1.722
	(3.16)			(−5.37)		(3.53)	(−3.07)	(−2.89)	(0.000)	
5	22.16	−27.82	45.31						80.78	1.610
	(4.87)	(−5.14)	(5.40)						(0.000)	
6	15.60			−17.60	620.16				75.64	1.776
	(3.32)			(−3.51)	(4.36)				(0.000)	
7	11.67	−24.64	46.32			6.95			72.18	1.525
	(2.14)	(−4.32)	(5.17)			(3.29)			(0.000)	
8	−3.22			−11.12	721.45	11.81			64.28	1.689
	(−0.51)			(−1.99)	(4.56)	(3.67)			(0.000)	
9	24.18	−29.87	46.76				−0.87	−3.45	77.72	1.579
	(4.79)	(−5.10)	(5.43)				(−0.26)	(−1.30)	(0.000)	
10	25.58			−28.87	863.71		−17.40	−11.20	50.45	1.656
	(3.52)			(−3.71)	(4.48)		(−3.29)	(−2.50)	(0.000)	
11	13.45	−31.48	42.37			12.82	−12.92	−1.89	68.14	1.434
	(2.19)	(−4.52)	(4.56)			(3.64)	(−3.03)	(−0.56)	(0.000)	
12	11.76			−28.75	755.61	13.76	−22.97	−5.65	46.98	1.573
	(1.40)			(−3.79)	(4.18)	(2.85)	(−3.99)	(−1.14)	(0.000)	
13	−5.65					6.48			108.01	1.892
	(−3.86)					(4.48)			(0.000)	
14	1.00						2.34	−3.00	109.55	1.950
	(21.47)						(1.05)	(−1.68)	(0.000)	
15	−16.64					17.18	−20.20	−2.66	92.84	1.880
	(−5.67)					(5.99)	(−6.25)	(−1.22)	(0.000)	

Table CIII

Test Results Based on the First- and Second-Stage GMM Estimates

This table presents results based on the first- and second-stage GMM estimates, for five pricing kernels in each case. It is otherwise identical to the corresponding part of Table IV.

Panel A: First-stage Results

Row	Const.	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	HJ-Distance
1	25.88	−24.95	8.57				−2.45	−10.67	1.819
	(4.22)	(−3.88)	(1.05)				(−0.56)	(−2.63)	
2	25.48			−27.05	246.38		−0.33	−16.02	1.685
	(3.83)			(−3.95)	(3.14)		(−0.05)	(−3.12)	
3	10.07	−27.65	13.85			17.21	−14.48	−4.01	1.693
	(1.25)	(−2.98)	(1.09)			(3.27)	(−2.21)	(−0.85)	
4	9.92			−28.57	277.29	16.21	−12.62	−11.12	1.553
	(1.12)			(−3.68)	(3.24)	(2.88)	(−1.49)	(−1.82)	
5	−12.12					12.76	−6.66	−1.73	1.880
	(−4.21)					(4.55)	(−1.97)	(−0.78)	

Panel B: Second-Stage Results

Row	Const.	mcg_i	vcg_i	mcg_{ae}	vcg_{ae}	R_M	SMB	HML	*J-Test*
1	28.71	−27.29	6.49				−0.31	−9.57	86.39
	(6.57)	(−6.26)	(1.22)				(−0.09)	(−2.85)	(0.000)
2	28.23			−30.82	327.80		−1.07	−19.55	53.03
	(4.55)			(−4.86)	(5.10)		(−0.21)	(−4.40)	(0.000)
3	18.05	−32.29	12.79			14.44	−10.55	−9.06	79.46
	(3.36)	(−5.20)	(1.72)			(3.23)	(−2.27)	(−2.53)	(0.000)
4	15.03			−31.85	312.83	13.97	−16.65	−12.80	45.08
	(1.96)			(−4.78)	(4.83)	(2.83)	(−2.32)	(−2.61)	(0.001)
5	−11.30					11.93	−9.89	−2.14	103.59
	(−4.63)					(5.01)	(−3.45)	(−1.12)	(0.000)

REFERENCES

Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review Papers and Proceedings* 80, 38–42.

Altonji, Joseph G., 1986, Intertemporal substitution in labor supply: Evidence from micro data, *Journal of Political Economy* 94, 176–215.

Altonji, Joseph G., and Aloysius Siow, 1987, Testing the response of consumption to income changes with (noisy) panel data, *Quarterly Journal of Economics* 102, 293–328.

Attanasio, Orazio P., and Guglielmo Weber, 1995, Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey, *Journal of Political Economy* 103, 1121–1157.

Balduzzi, Pierluigi, and Tong Yao, 2000, Does heterogeneity matter for asset pricing? Working paper, Boston College.

Basak, Suleyman, and Domenico Cuoco, 1998, An equilibrium model with restricted stock market participation, *Review of Financial Studies* 11, 309–341.

Black, Fischer, Michael Jensen, and Myron Scholes, 1972, The capital asset pricing model: Some empirical tests, in Michael Jensen, ed. *Studies in the Theory of Capital Markets* (Praeger, New York, NY).

Brav, Alon, George Constantinides, and Christopher Geczy, 2002, Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, *Journal of Political Economy* 110, 793–824.

Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investing opportunities, *Journal of Financial Economics* 7, 265–296.

Breeden, Douglas T., Michael Gibbons, and Robert Litzenberger, 1989, Empirical tests of the consumption-based CAPM, *Journal of Finance* 44, 231–262.

Browning, Martin, Angus S. Deaton, and Margaret Irish, 1985, A profitable approach to labor supply and commodity demands over the life-cycle, *Econometrica* 53, 503–544.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behaviour, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., and John H. Cochrane, 2000, Explaining the poor performance of consumption-based asset pricing models, *Journal of Finance* 55, 2863–2878.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).

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

Cochrane, John H., 1996, A cross sectional test of an investment-based asset pricing model, *Journal of Political Economy* 104, 572–621.

Cochrane, John H., 2001, *Asset Pricing* (Princeton University Press, Princeton, NJ).

Cochrane, John H., and Lars Peter Hansen, 1992, Asset pricing lessons for macroeconomics, in Oliver Blanchard and Stanley Fischer, eds. *1992 NBER Macroeconomics Annual* (National Bureau of Economic Research, Cambridge, MA).

Cogley, Timothy, 2002, Idiosyncratic risk and the equity premium: Evidence from the consumer expenditure survey, *Journal of Monetary Economics* 49, 309–334.

Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.

Constantinides, George M., 2002, Rational asset prices, *Journal of Finance* 57, 1567–1591.

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

Constantinides, George M., John B. Donaldson, and Rajnish Mehra, 2002, Junior can't borrow: A new perspective on the equity premium puzzle, *Quarterly Journal of Economics* 117, 269–296.

Detemple, Jerome, and Fernando Zapatero, 1991, Asset prices in an exchange economy with habit formation, *Econometrica* 59, 1633–1657.

Dumas, Bernard, 1989, Two-person dynamic equilibrium in the capital market, *Review of Financial Studies* 2, 157–188.

Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of asset returns: An empirical analysis, *Journal of Political Economy* 99, 263–287.

Fama, Eugene F., and James MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 71, 607–636.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

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

Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.

Fama, Eugene F., and G. William Schwert, 1977, Human capital and capital market equilibrium, *Journal of Financial Economics* 4, 115–146.

Ferson, Wayne E., 1995, Theory and empirical testing of asset pricing models, in R. A. Jarrow, V. Maksimovic and W. T. Ziemba, eds. *Handbooks in Operations Research and Management Science* (Elsevier Science B.V, Amsterdam).

Ferson, Wayne E., and George M. Constantinides, 1991, Habit persistence and durability in aggregate consumption: Empirical tests, *Journal of Financial Economics* 29, 199–240.

Ferson, Wayne E., and Stephen Foerster, 1994, Finite sample properties of the generalized method of moments tests of conditional asset pricing models, *Journal of Financial Economics* 36, 29–55.

Ferson, Wayne E., and Campbell R. Harvey, 1992, Seasonality and consumption-based asset pricing, *Journal of Finance* 47, 511–552.

Grossman, Sanford J., Angelo Melino, and Robert J. Shiller, 1987, Estimating the continuous time consumption-based asset pricing model, *Journal of Business and Economics Statistics* 5, 315–327.

Hansen, Lars Peter, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.

Hansen, Lars Peter, and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars Peter, and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557–590.

Hansen, Lars Peter, and Kenneth Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

Hansen, Lars Peter, and Kenneth Singleton, 1983, Stochastic consumption, risk aversion and the temporal behavior of asset returns, *Journal of Political Economy* 91, 249–265.

Hansen, Lars Peter, and Kenneth Singleton, 1984, Errata, *Econometrica* 52, 267–268.

Hayashi, Fumio, 1985, The permanent income hypothesis and consumption durability: Analysis based on Japanese panel data, *Quarterly Journal of Economics* 100, 1083–1113.

Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff, 1994, Risk sharing across and within families, *Econometrica* 64, 261–294.

Heaton, John, 1993, The interaction between time-nonseparable preferences and time aggregation, *Econometrica* 61, 353–385.

Heaton, John, 1995, An empirical investigation of asset pricing with temporally dependent preference specifications, *Econometrica* 63, 681–717.

Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets for risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.

Jacobs, Kris, 1999, Incomplete markets and security prices: Do asset-pricing puzzles result from aggregation problems? *Journal of Finance* 54, 123–164.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–53.

Kan, Raymond, and Chu Zhang, 1999, Two-pass tests of asset pricing models with useless factors, *Journal of Finance* 54, 203–235.

Kocherlakota, Narayana, 1996, The equity premium: It's still a puzzle, *Journal of Economic Literature* 34, 42–71.

Lettau, Martin, and Sydney Ludvigson, 2002, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.

Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Mace, Barbara, 1991, Full insurance in the presence of aggregate uncertainty, *Journal of Political Economy* 99, 928–956.

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

Mankiw, N. Gregory, Julio Rotemberg, and Lawrence Summers, 1985, Intertemporal substitution in macroeconomics, *Quarterly Journal of Economics* 100, 225–252.

Mankiw, N. Gregory, and Matthew Shapiro, 1986, Risk and return: Consumption beta versus market beta, *Review of Economics and Statistics* 68, 452–459.

Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29, 113–135.

Mayers, David, 1972, Nonmarketable assets in capital market equilibrium, in Michael Jensen, ed. *Studies in the Theory of Capital Markets* (Praeger Publishers Inc., New York, NY).

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Miron, Jeffrey A., 1986, Seasonal fluctuations and the life cycle-permanent income model of consumption, *Journal of Political Economy* 94, 1258–1279.

Roll, Richard, 1977, A critique of asset pricing theory's tests: Part 1, *Journal of Financial Economics* 4, 129–176.

Sarkissian, Sergei, 2003, Incomplete consumption risk sharing and currency risk premiums, *Review of Financial Studies* 16, 983–1005.

Sundaresan, Suresh, 1989, Intertemporally dependent preferences and the volatility of consumption and wealth, *Review of Financial Studies* 2, 73–88.

Telmer, Chris I., 1993, Asset pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.

Telmer, Chris I., Kjetil Storesletten, and Amir Yaron, 1997, Asset pricing with idiosyncratic risk and overlapping generations, GSIA Working paper no. 1997-E226, Carnegie Mellon University.

Vissing-Jorgensen, Annette, 2002, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of Political Economy* 110, 825–853.

Zeldes, Stephen P., 1989, Consumption and liquidity constraints, *Journal of Political Economy* 97, 305–346.