Explaining the Poor Performance of Consumption-Based Asset Pricing Models

Citation

Campbell, John Y., and John H. Cochrane. 2000. Explaining the poor performance of consumption-based asset pricing models. *Journal of Finance* 55(6): 2863-2878.

Published Version

<http://dx.doi.org/10.1111/0022-1082.00310>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3163265>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

NBER WORKING PAPER SERIES

EXPLAINING THE POOR PERFORMANCE
OF CONSUMPTION-BASED
ASSET PRICING MODELS

John Y. Campbell
John H. Cochrane

Working Paper 7237
<http://www.nber.org/papers/w7237>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 1999

Campbell's research is supported by the National Science Foundation via a grant administered to the NBER.; Cochrane's research is supported by the National Science Foundation via a grant administered to the NBER and by the Graduate School of Business. We thank Andrea Eisfeldt, Andrew Abel, George Constantinides, Lars Hansen, John Heaton, Robert Lucas, and an anonymous referee for helpful comments. All opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

© 1999 by John Y. Campbell and John H. Cochrane. All rights reserved. Short sections of text, not to

exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Explaining the Poor Performance of
Consumption-Based Asset Pricing Models

John Y. Campbell and John H. Cochrane

NBER Working Paper No. 7237

July 1999

JEL No. G00

ABSTRACT

The poor performance of consumption-based asset pricing models relative to traditional portfolio-based asset pricing models is one of the great disappointments of the empirical asset pricing literature. We show that the external habit-formation model economy of Campbell and Cochrane (1999) can explain this puzzle. Though artificial data from that economy conform to *a* consumption-based model by construction, the CAPM and its extensions are much better approximate models than is the standard power utility specification of the consumption-based model. Conditioning information is the central reason for this result. The model economy has one shock, so when returns are measured at sufficiently high frequency the consumption-based model and the CAPM are equivalent and perfect *conditional* asset pricing models. However, the model economy also produces time-varying expected returns, tracked by the dividend-price ratio. Portfolio-based models capture some of this variation in state variables, which a state-independent function of consumption cannot capture, and so portfolio-based models are better approximate *unconditional* asset pricing models.

John Y. Campbell
Department of Economics
Harvard University
Littauer Center 213
Cambridge, MA 02138
and NBER
john_campbell@harvard.edu

John H. Cochrane
Graduate School of Business
University of Chicago
1101 E. 58th. St.
Chicago, IL 60637
and NBER
john.cochrane@gsb.uchicago.edu

1. Introduction

The development of consumption-based asset pricing theory ranks as one of the major advances in financial economics during the last two decades. The classic papers of Lucas (1978), Breeden (1979), Grossman and Shiller (1981), and Hansen and Singleton (1982, 1983) showed how a simple relation between consumption and asset returns captures the implications of complex dynamic intertemporal multifactor asset pricing models.

Unfortunately, consumption-based asset pricing models have proved disappointing empirically. Hansen and Singleton (1982, 1983) formulated a canonical consumption-based model in which a representative investor has time-separable power utility of consumption. They rejected the model on US data, finding that it could not simultaneously explain the time-variation of interest rates and the cross-sectional pattern of average returns on stocks and bonds. Wheatley (1988) rejected the model on international data.

All models can be rejected, and the more important issue is which approximate models are most useful. Alas, the canonical consumption-based model has performed no better, and in many respects worse, than even the simple static CAPM. Mankiw and Shapiro (1986) regress the average returns of the 464 NYSE stocks that were continuously traded from 1959 to 1982 on their market betas, on consumption growth betas, and on both betas. They find that market betas are more strongly and robustly associated with the cross-section of average returns, and they find that market betas drive out consumption betas in multiple regressions. Breeden, Gibbons and Litzenberger (1989) study industry and bond portfolios, finding roughly comparable performance of the CAPM and a model that uses a mimicking portfolio for consumption growth as the single factor, after adjusting the consumption-based model for measurement problems in consumption. Cochrane (1996) finds that the traditional CAPM substantially outperforms the canonical consumption-based model in pricing size portfolios. For example, he reports a root mean square pricing error (alpha) of 0.094 percent per quarter for the CAPM and 0.54 percent per quarter for the consumption-based model.

More recently, multifactor models have improved on the CAPM. Jagannathan and Wang (1996) and Cochrane (1996) extend the traditional CAPM by scaling the market factor with “price ratio” variables that reveal market expectations, such as the dividend-price ratio or the term premium. This extended CAPM can be interpreted as a conditional CAPM, or as an unconditional multifactor model. Cochrane (1996) reports pricing errors about half those of the static CAPM on size portfolios. Chen, Roll, and Ross (1986) and Jagannathan and Wang (1996) reduce pricing errors by adding macroeconomic factors, and Fama and French (1993) use size and book-market factors to dramatically reduce the CAPM’s pricing errors on size and book-market

sorted portfolios.

The canonical consumption-based model has failed perhaps the most important test of all, the test of time. 25 years after the development of the consumption-based model, almost all applied work in finance still uses portfolio-based models to correct for risk, to digest anomalies, to produce cost of capital estimates, and so forth.

This history is often interpreted as evidence against consumption-based models in general rather than against particular utility functions, particular specifications of temporal nonseparabilities such as habit persistence or durability, and particular choices of consumption data and data-handling procedures. But this conclusion is internally inconsistent, because all current asset pricing models are derived as specializations of the consumption-based model rather than as alternatives to it. All current models predict that expected returns should line up against covariances of returns with some function of consumption (possibly including leads and lags). For example, the CAPM is derived by specializing the consumption-based model to two periods, quadratic time-separable utility, and no labor income (or to log utility and lognormally distributed returns; or to quadratic utility and i.i.d. returns; see Cochrane 1999 for textbook derivations). Portfolio-based models are *not* derived by the assumption of explicit frictions that de-link consumption from asset returns. One cannot believe that the CAPM does hold, but consumption-based models, as a class, fundamentally do not.

Still, the canonical consumption-based model does poorly in practice relative to factor-pricing models that use portfolio returns as risk factors, and it is important to understand why this is so. In this paper, we examine this issue using artificial data from the Campbell and Cochrane (1999) model economy. A consumption-based model does hold, exactly and by construction, yet we find that the CAPM outperforms the canonical specification of the consumption-based model, and that a multifactor extension of the CAPM performs better still.

Since we study artificial data from a fully specified economy, we are able to analyze the economic reasons for these results. Conditioning information is the central element of the story. The model has only one shock, so as the measurement interval shrinks (eliminating the effects of nonlinearities), consumption growth and the market return are both perfectly *conditionally* correlated with the stochastic discount factor. Thus, consumption growth or the market return both provide a perfect *conditional* asset pricing model; conditional expected returns line up perfectly with conditional betas on the market portfolio or conditional betas on consumption growth.

However, returns are not i.i.d., as the model economy generates time-varying expected returns which can be forecast by dividend-price ratios. This means that *unconditional* correlations need not match conditional correlations. It turns out that the market return is better *unconditionally* correlated with the true discount factor

than is consumption growth and thus the market return is a better proxy for an unconditional asset pricing model. The reason for this is that the market return is affected when the price-dividend ratio changes. The market return therefore reflects variation in this state variable that consumption growth does not. Equivalently, the stochastic discount factor is a state-dependent function of consumption growth; the market return captures some of this state-dependence as well as some correlation with consumption growth shocks.

One can always argue in principle that perhaps the utility function is misspecified, but it has been hard to believe that plausible changes in utility functions could explain the amount by which portfolio-based models outperform the canonical consumption-based model with power utility. Our contribution is to show in an explicit quantitative example that, in fact, portfolio-based models can outperform the canonical consumption-based model by the amount we see in the data, even when a slightly more complex consumption-based model holds by construction.

2. The Economic Model

To generate time-varying expected returns, the model economy adds habit persistence to the standard consumption-based specification. As bad shocks drive consumption down towards the habit level, risk aversion rises, stock prices decline, and expected returns rise. Campbell and Cochrane (1999) describe the model in detail, and motivate the ingredients.

Consumption growth is an i.i.d. lognormal endowment process,

$$\Delta c_{t+1} = g + v_{t+1}; \quad v_{t+1} \sim i.i.d. \mathcal{N}(0, \sigma^2). \quad (1)$$

Identical agents maximize the utility function

$$E \sum_{t=0}^{\infty} \delta^t \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma}. \quad (2)$$

Here X_t is the level of habit, and δ is the subjective or time discount factor. Lowercase letters denote logarithms of uppercase letters, $c_t = \ln C_t$, etc.

It is convenient to capture the relation between consumption and habit by the *surplus consumption ratio*

$$S_t = \frac{C_t - X_t}{C_t}.$$

A process for the surplus consumption ratio specifies how habit X_t responds to the history of consumption. The log surplus consumption ratio evolves as

$$s_{t+1} = (1 - \phi)\bar{s} + \phi s_t + \lambda(s_t)(c_{t+1} - c_t - g). \quad (3)$$

ϕ , g and $\bar{s}$ are parameters. It is convenient to specify that the parameter g in (3) is equal to the mean consumption growth rate g , but this is not essential.

The *sensitivity function* $\lambda(s_t)$ in (3) controls the sensitivity of s_{t+1} and thus habit x_{t+1} to contemporaneous consumption c_{t+1} . It is given by

$$\lambda(s_t) = \begin{cases} \frac{1}{\bar{s}} \sqrt{1 - 2(s_t - \bar{s})} - 1, & s_t \leq s_{\max} \\ 0 & s_t \geq s_{\max} \end{cases} \quad (4)$$

where

$$\bar{s} \equiv \sigma \sqrt{\frac{\gamma}{1 - \phi}}, \quad (5)$$

and $s_{\max}$ is the value of s_t at which the square root in (4) runs into zero,

$$s_{\max} \equiv \bar{s} + \frac{1}{2} (1 - \bar{s}^2). \quad (6)$$

The specification is not as complex as it seems at first glance: Equations (3) and (4) are almost a familiar square-root process. This specification of $\lambda(s_t)$ achieves three objectives: 1) The riskfree interest rate is constant. As consumption declines toward habit, consumers in a nonstochastic economy would want to borrow, driving up interest rates. However, as consumption declines, $\lambda(s_t)$ rises. This rise acts in (3) like an increase in risk, which increases precautionary savings, lowering interest rates. Our specification of $\lambda(s_t)$ makes these two effects offset exactly. 2) Habit is predetermined at the steady state $s_t = \bar{s}$. 3) Habit is also predetermined near the steady state, or, equivalently, habit moves non-negatively with consumption everywhere. It also turns out that $\lambda(s_t)$ must rise as s_t falls in order to generate a time-varying conditional Sharpe ratio.

Marginal utility and asset prices

We assume that habit is *external*; people want to “keep up with the Joneses” as in Abel (1990). Then, marginal utility is given by

$$u_c(C_t, X_t) = (C_t - X_t)^{-\gamma} = S_t^{-\gamma} C_t^{-\gamma}.$$

The intertemporal marginal rate of substitution, or stochastic discount factor, is

$$M_{t+1} \equiv \delta \frac{u_c(C_{t+1}, X_{t+1})}{u_c(C_t, X_t)} = \delta \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma}. \quad (7)$$

The log marginal rate of substitution is

$$\begin{aligned} m_{t+1} &= \ln(\delta) - \gamma (\Delta s_{t+1} + \Delta c_{t+1}) \\ &= \ln(\delta) - \gamma g - \gamma(\phi - 1)(s_t - \bar{s}) - \gamma(1 + \lambda(s_t))v_{t+1}. \end{aligned} \quad (8)$$

This variable is conditionally normally distributed. The external habit specification is convenient, because it allows us to ignore terms by which current consumption might affect future habits. In Campbell and Cochrane (1999) we argue that many of the aggregate properties of the model are substantially unaffected by the choice of an external rather than an internal specification.

The real riskfree interest rate is the reciprocal of the conditionally expected stochastic discount factor,

$$R_t^f = 1/E_t(M_{t+1}).$$

Using equation (8) and taking the expectation of the lognormal random variable M , the log riskfree rate is

$$r_t^f = -\ln(\delta) + \gamma g - \frac{1}{2}\gamma(1 - \phi). \quad (9)$$

We use a claim to the consumption stream to model the market portfolio. Its price-dividend ratio satisfies

$$\frac{P_t}{C_t}(s_t) = E_t \left[M_{t+1} \frac{C_{t+1}}{C_t} \left(1 + \frac{P_{t+1}}{C_{t+1}}(s_{t+1}) \right) \right]. \quad (10)$$

We solve this functional equation numerically on a grid for the state variable s_t , using numerical integration and interpolation of the $P/C(s)$ function to evaluate the conditional expectation. Given the price-consumption ratio as a function of state and the state transition equation (3), we can simulate returns and other interesting quantities.

We also model a claim to dividends that are imperfectly correlated with consumption. We specify that log dividend growth is also i.i.d., and has correlation coefficient ρ with aggregate consumption growth,

$$\Delta d_{t+1} = g + w_{t+1}; \quad w_{t+1} \sim i.i.d. \mathcal{N}(0, \sigma_w^2), \quad corr(w_t, v_t) = \rho. \quad (11)$$

The price-dividend ratio of this claim is also a function of the state variable s_t , and we find it in the same way.

We simulate our model at a monthly frequency. We construct time-averaged annual consumption data by summing consumption during the year. This procedure is a crude way to capture the effect of time-aggregation in measured consumption.

Choosing Parameters

We use the same parameter values as in Campbell and Cochrane (1999), calibrated to postwar (1947-1995) annual NIPA nondurable and services per capita consumption together with data from the CRSP value-weighted NYSE stock portfolio. Table 1

summarizes the parameter choices. The mean and standard deviation of log consumption growth, g and σ , match the consumption data. We choose the serial correlation parameter ϕ to match the serial correlation of the log price-dividend ratio. We choose the riskfree rate to match the average real return on Treasury bills. We choose the utility curvature parameter γ to match the market Sharpe ratio. We calibrate the standard deviation and consumption growth correlation of the dividend process from the CRSP value-weighted return data as well. Parameters δ , $\bar{S}$, $S_{\max}$ follow from these choices via (9), (5) and (6).

Campbell and Cochrane (1999) show that the model with these parameter choices matches a wide variety of phenomena including the equity premium, the predictability of stock returns from price-dividend ratios, violations of volatility tests and the leverage effect by which lower prices imply more volatile returns. This point is important for the current exercise: Our story is based on time-varying conditioning information, and one wants reassurance that the assumed time-variation in return distributions is sensible.

Table 1. Parameter choices.

Assumed parameters	
Mean consumption growth* g , %	1.89
Standard deviation of consumption growth* σ , %	1.50
Log riskfree rate* r^f , %	0.94
Persistence coefficient* ϕ	0.87
Utility curvature γ	2.00
Standard deviation of dividend growth* σ_w , %	11.2
Correlation between consumption, dividend growth ρ	0.2
Implied parameters	
Subjective discount factor* δ	0.89
Steady state surplus consumption ratio $\bar{S}$	0.057
Maximum surplus consumption ratio $S_{\max}$	0.094

Starred (*) table entries are annualized values, e.g. $12g$, $\sqrt{12}\sigma$, $12r^f$, ϕ^{12} , δ^{12} , $\sqrt{12}\sigma_w$ since the model is simulated at a monthly frequency. δ , $\bar{S}$ and $S_{\max}$ are calculated using (9), (5) and (6) respectively.

3. Implications for Cross-Sectional Tests of Asset Pricing Models

Now we can answer our basic question. How do the standard power utility consumption-based model, the CAPM, and multifactor extensions compare in artificial data from our model?

False Models

We specify the alternative asset pricing models in terms of their stochastic discount factors, which we denote Y . We use Y to distinguish false discount factor proxies from the true discount factor M . The alternative models are:

1. The canonical consumption-based model with time-separable power utility,

$$Y_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\eta}. \quad (12)$$

We use β and η to emphasize that these coefficients need not be equal to the parameters δ and γ of the data-generating model. This is the classic form studied by Hansen and Singleton (1982).

2. A consumption-based linear factor model,

$$Y_{t+1} = a + b \left(\frac{C_{t+1}}{C_t} \right). \quad (13)$$

A discount factor that is a linear function of a set of variables implies that expected returns are linear in betas, i.e., multiple regression coefficients of returns on those variables. Thus (13) corresponds to tests that compare expected returns with consumption betas, without imposing the nonlinear structure of (12), as in Mankiw and Shapiro (1986) and Breeden, Gibbons and Litzenberger (1989).

3. The traditional static CAPM,

$$Y_{t+1} = a + bR_{t+1}^w, \quad (14)$$

where R^w is the consumption claim (wealth) return.

4. A multifactor or conditional CAPM,

$$Y_{t+1} = [a_0 + a_1(d_t - p_t)] + [b_1 + b_2(d_t - p_t)] R_{t+1}^w \quad (15)$$

$$= a_0 + a_1(d_t - p_t) + b_1 R_{t+1}^w + b_2 [(d_t - p_t) R_{t+1}^w]. \quad (16)$$

The first equation expresses this model as a conditional CAPM—a CAPM with time-varying coefficients. The second equation expresses the same model as an unconditional multifactor asset pricing model.

Our model is particularly good motivation for this form of a conditional CAPM. The canonical consumption-based model with power utility implies

$$y_{t+1} = \ln(\beta) - \eta(c_{t+1} - c_t). \quad (17)$$

The true model is, from (8),

$$m_{t+1} = \ln(\delta) + [\gamma g \lambda(s_t) - \gamma(\phi - 1)(s_t - \bar{s})] - \gamma(1 + \lambda(s_t))(c_{t+1} - c_t). \quad (18)$$

The true model has the same form as the canonical consumption-based model (17), except that it makes the intercept and slope on consumption growth functions of the slow-moving state variable s_t . The multifactor model (15) also has this form, but it is driven by the market return rather than by consumption growth.

Pricing errors

We want to know, how big are the pricing errors of the false models when applied to a cross-section of assets? Our single-shock model does not naturally give rise to an interesting cross-section of assets such as the size, book-to-market, industry, government bond, corporate bond, or international portfolios studied in the asset pricing literature. Therefore, we use a distance measure introduced by Hansen and Jagannathan (1997) and related to Shanken (1987) to find the maximum pricing error that the false models can produce.

Hansen and Jagannathan show that the maximum possible pricing error, expressed in Sharpe ratio units as expected return error (Jensen's alpha) per unit of standard deviation, is proportional to the standard deviation of the difference between the true and false discount factors. To express these ideas formally, let $E^Y(\xi)$ denote the expected value of a payoff ξ predicted by the false discount factor Y . Then

$$\max_{\{\xi\}} \frac{|E^Y(\xi) - E(\xi)|}{\sigma(\xi)} = \frac{\sigma(M - Y)}{E(M)}. \quad (19)$$

The left hand side of (19) is the definition of the maximum pricing error per unit of standard deviation. The right hand side, the Hansen-Jagannathan distance measure, relates the pricing error to the standard deviation of the difference between true and false discount factors.

Each of our false discount factor models has free parameters. We fix the free parameters in two alternative ways, either by minimizing the maximum pricing error, or by "estimating" parameters that best price the riskfree rate and the market return. When the discount factor Y is a linear function of factors, $Y = \mathbf{b}'\mathbf{f}$, these procedures are related. In this case, minimizing the maximum pricing error is the same as ensuring that the false model correctly prices the factors,

$$\min_{\{\mathbf{b}\}} E[(M - \mathbf{b}'\mathbf{f})^2] \Rightarrow E(M\mathbf{f}) = E(\mathbf{f}(\mathbf{f}'\mathbf{b})) = E(Y\mathbf{f}). \quad (20)$$

Minimizing the maximum pricing error is also equivalent to an OLS regression of the true discount factor M on the factors $\mathbf{f}$. From (20),

$$\mathbf{b} = E(\mathbf{f}\mathbf{f}')^{-1}E(M\mathbf{f}).$$

Derivation and further interpretation

To understand the Hansen-Jagannathan result, consider a payoff ξ with price P . The pricing relation $P = E(M\xi)$ implies

$$E(\xi) = \frac{P}{E(M)} - \frac{\text{cov}(M, \xi)}{E(M)}.$$

The expected payoff (return) predicted by the false discount factor Y is given by

$$E^Y(\xi) \equiv \frac{P}{E(Y)} - \frac{\text{cov}(Y, \xi)}{E(Y)}.$$

Suppose the approximate model gives the same average price of a riskfree rate, i.e. $E(Y) = E(M)$. We pick parameters to ensure this equality in our application, in order to focus entirely on models' ability to correct for risk. Then the expected return error is

$$|E^Y(\xi) - E(\xi)| = \left| \frac{\text{cov}(M - Y, \xi)}{E(M)} \right| \leq \frac{\sigma(M - Y)\sigma(\xi)}{E(M)}.$$

The payoff $\xi^* = M - Y$ makes the inequality tight. Sensibly, this worst-priced payoff is perfectly correlated with the difference between the true and false discount factors. Hansen and Jagannathan (1997) derive the result in a much more general setting.

It is both the advantage and disadvantage of this technique that it depends only on the model, not on the set of test portfolios. Approximate models can work well on some portfolios but poorly on others. The CAPM, for example, works well on beta-sorted stock portfolios, decently on industry- and size-sorted portfolios, but poorly on portfolios sorted by book-market ratio. Kandel and Stambaugh (1995) and Roll and Ross (1994) show how the pricing errors of an approximate model can depend dramatically on arbitrary choices of test portfolios.

The advantage, then, is that the Hansen-Jagannathan procedure eliminates this arbitrariness by evaluating the pricing error of the worst possible portfolio, the one generates the largest possible pricing error. The search for the worst-priced payoff extends over all possible contingent claims, including all dynamic strategies. $\xi^* = M - Y$ will typically be a function of consumption growth as well as asset returns.

This is also the disadvantage, as our experience with the relative performance of the CAPM and consumption-based model is based on a quite limited set of assets, especially when compared to the set of all contingent claims. For example, if the maximum pricing error of a false model occurs for a portfolio that is a highly non-linear function of consumption growth, does that really tell us much about which models price stock portfolios well? Even within the limited set of assets that has been examined, results seem to be sensitive to the asset choice; Breeden, Gibbons and Litzenberger (1989) used industry and bond portfolios, and found better results for

the canonical consumption-based model than did Mankiw and Shapiro (1986) using individual stocks or Cochrane (1996) using size portfolios.

With a specific set of traded assets in mind, one could generalize the Hansen-Jagannathan technique to characterize only the pricing errors of traded assets. Equation (19) generalizes to

$$\max_{\{\xi \in \mathbf{X}\}} \frac{|E^Y(\xi) - E(\xi)|}{\sigma(\xi)} = \min_{\{M: P=E(MX) \forall X \in \mathbf{X}\}} \frac{\sigma(M - Y)}{E(M)} \quad (21)$$

where $\mathbf{X}$ denotes the space of traded assets, and P and X denote a price and payoff of a traded asset.

One could generate a limited set of traded assets, by generating multiple dividends of the form (11) and pricing them via the true model. One could then evaluate pricing errors of false models via (21). However, the fundamental economic characteristics that drive the cross-sectional variation in observed equity returns are poorly understood and are unlikely to be well modeled by simple processes such as (11). Therefore, we do not limit the space of portfolios over which to search for large pricing errors.

Results

Table 2 gives the maximum pricing errors in artificial data from the Campbell-Cochrane model. These are population moments, recovered from a simulation of 100,000 months of artificial data time-aggregated to an annual frequency.

We start with the static CAPM, using the consumption-claim return as the market return, in row a. The maximum pricing errors have a 0.40 Sharpe ratio, or 7.9% average return at a 20% standard deviation. This is roughly the size of the worst CAPM pricing errors in the literature. For example, Fama and French (1993) find the CAPM does nothing to explain the roughly 10% expected return variation across book-market sorted portfolios, and they report that the high-minus-low book-market portfolio earns a Sharpe ratio roughly that of the market portfolio, despite a very low market beta.

The scaled CAPM in row b does a bit better than the static CAPM. The improvement in performance is not dramatic, possibly because we search for the largest pricing error among all contingent claims, rather than among a set of portfolios sorted on the same basis as the factors. Since these are population moments, there is no mechanical improvement as one adds additional factors.

Table 2. Maximum pricing errors from approximate asset pricing models

Panel A: Basic Results					
Model		α/σ	$\alpha(\%)$	$\rho_{Y,M}$	η
a)	CAPM, $Y_{t+1} = a + bR_{t+1}^w$	0.40	7.9	0.77	
b)	Scaled CAPM, $Y_{t+1} = a_0 + a_1(pd_t) + [b_0 + b_1(pd_t)]R_{t+1}$	0.36	7.1	0.82	
c)	Power utility, $Y_{t+1} = \beta(C_{t+1}/C_t)^{-\eta}$	0.52	10.3	0.56	29
d)	Power utility, β, η chosen to price R^w, R^f	1.01	20.2	0.56	78
e)	Risk-neutral, $Y_{t+1} = 1/R^f$	0.62	12.5	0	

Panel B: Variations					
Model		α/σ	$\alpha(\%)$	$\rho_{Y,M}$	
f)	CAPM, monthly simulated data	0.13	2.5	0.97	
g)	Power utility model, monthly simulated data	0.23	4.7	0.91	
h)	Consumption factor model, $Y_{t+1} = a + b(C_{t+1}/C_t)$	0.54	10.8	0.50	
i)	Consumption factor model, a, b chosen to price R^w, R^f	0.93	18.5	0.50	
j)	CAPM, dividend claim return $Y_{t+1} = a + bR_{t+1}^d$	0.48	9.5	0.65	
k)	Scaled CAPM, dividend claim return	0.35	7.0	0.83	

Note to Table 2. α/σ gives the maximum expected return pricing error per unit of standard deviation produced by each of the false discount factor models. This is calculated by $\alpha/\sigma = \sigma(M - Y)/E(M)$. $\alpha(\%)$ evaluates the pricing error at a 20% annual standard deviation, expressed as an annual percentage average return. η gives the estimated curvature coefficient in the consumption-based models. $\rho_{Y,M}$ gives the unconditional correlation between true and false discount factors. The monthly results in rows *f* and *g* are annualized by multiplying by $\sqrt{12}$.

We always choose parameters so that the false discount factor model correctly prices the riskfree rate. In rows a,b,c we minimize the maximum pricing error, which is equivalent to correctly pricing the factors in row a,b. In row d we pick η to correctly price the consumption claim return. Results are based on 100,000 months of simulated data, which are time-aggregated to an annual frequency except where otherwise noted.

Row c presents the canonical consumption-based model with power utility. In this row, we estimated the parameters of the model to minimize the maximum pricing error. This is as good as the power-utility model can do, and it fares a good deal worse than the CAPM. The maximum pricing error is about 30% larger than that of the CAPM. We estimate a large coefficient of risk aversion ($\eta = 29$), as in real data.

The performance of the power-utility model is sensitive to parameter choices. Choosing parameters by minimizing the maximum pricing error forces the CAPM to correctly price the market proxy, but this is not true for the consumption-based model. In row d we pick η to correctly price the market return. (Precisely, we minimize the pricing error, and it turns out that $\eta = 78$ sets the error to zero. We always choose β to match the risk-free rate.) This is the same condition used to pick CAPM parameters, and it more closely mirrors practice, where we pick parameters to minimize pricing errors of a cross-section of traded assets rather than to minimize the maximum pricing error.

The pricing errors in row d are now more than double those of the CAPM: a 1.01 Sharpe ratio corresponding to a 20.2% average return for a 20% standard deviation. This pricing error is larger than the spread in expected returns in most studies. It is also larger than the pricing error of a risk-neutral model shown in row e. The consumption-based model is literally worse than useless.

To understand these results, note that there is only one shock in the Campbell-Cochrane model economy, so consumption growth, returns, and the discount factor become perfectly *conditionally* correlated as the time interval of the model shrinks. (Nonlinearities are the only reason the correlation isn't perfect in the discrete-time version of the model.) However, the sensitivities of consumption growth, returns, and the discount factor to the underlying shock vary over time and from each other. Thus, consumption growth or returns are imperfectly *unconditionally* correlated with the discount factor, and are thus imperfect proxies for unconditional asset-pricing models.

The CAPM performs better because the stock return is *more closely* unconditionally correlated with the marginal rate of substitution than is consumption growth, as one can see from the correlation coefficients in Table 2. Recall that the marginal rate of substitution is given by

$$M_{t+1} = \delta \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma}. \quad (22)$$

Changes in the surplus consumption ratio reduce the unconditional correlation between consumption growth and marginal utility growth. The stock return moves when consumption (its dividend) moves, but the stock return also moves when there is a change in the surplus consumption ratio, so the stock return is better correlated with marginal utility growth. Time-aggregation (and, in real life, measurement error) add to the advantage of the stock return as a proxy.

The scaled CAPM controls for variation in conditioning information by making parameters functions of state. It produces somewhat lower pricing errors than the CAPM, and somewhat higher correlation with the true discount factor. A conditional

CAPM that uses a square root transformation of a state variable does even better, as one might expect from (8).

Panel B of Table 2 presents several variations on the theme. Rows f and g present CAPM and power-utility results in our monthly simulated data with no time-aggregation. The CAPM produces about half the pricing errors of the power-utility model. This calculation verifies that the relative performance of the two models is centrally due to conditioning information, and not just to time-aggregation in the annual consumption data. Variation in conditioning information is more important at shorter investment horizons, which is why the relative performance of the CAPM is even better in monthly artificial data, despite the lack of consumption time-aggregation.

Rows h and i present pricing errors for a model that uses consumption growth as a factor, but does not impose the nonlinear specification of power utility, as in Breeden, Gibbons and Litzenberger (1989) and Mankiw and Zeldes (1985). The pricing errors are almost identical to those of the nonlinear power-utility model. This calculation verifies that the poor performance of the consumption-based model is not due to the occasional spectacular outliers that result from raising consumption growth to the -78 power.

Rows j and k use a claim to dividends poorly correlated with consumption as the market proxy in the CAPM. We present these results with some reservations. First, it is true that empirical work almost always uses the stock market portfolio rather than a larger wealth portfolio in evaluating the CAPM. On the other hand, this empirical work also almost always only tests the CAPM on similar assets (stocks), rather than testing the model on bonds, options, foreign exchange, real estate, human capital or other returns. Since we do not limit the space of returns when looking for pricing errors, it seems to us that the consumption-claim CAPM is a more appropriate benchmark than the dividend-claim CAPM. Second, our model of dividend growth is unrealistic in an important dimension. The correlation between dividend and consumption growth undoubtedly rises at longer horizons. The two series are very likely to be cointegrated, as it is hard to envision a growth path in which the ratio of dividends to aggregate consumption rises to infinity or declines to zero. Our model of dividends was useful for showing that the aggregate asset pricing results in Campbell and Cochrane (1999) are insensitive to which aggregate asset *is priced*, but one would want a better dividend model to use the return as a factor to price other assets. Such a model would involve the addition of at least one additional state variable, the dividend-consumption ratio, which is not worth the space and complexity it would entail for the purposes of this robustness check.

At any rate, the static dividend-claim CAPM in row j has substantially higher pricing errors than the static consumption-claim CAPM in row a, yet still lower than those of the power-utility model. When we scale the dividend-claim CAPM in row

k, we find that its pricing errors fall close to those of the scaled consumption-claim CAPM in row b. This shows that scaling is more important when using the noisier dividend-claim return as a proxy. (In monthly simulated data, without consumption time-aggregation, the unscaled dividend-claim CAPM slightly underperforms the power-utility model.)

It is tempting to continue, using our artificial data to evaluate other models and to study the effects of alternative data transformations. One could examine the performance of more complex consumption-based models, such as Ferson and Constantinides' (1991) habit specification, or Epstein and Zin's (1991) non-expected-utility formulation. One could try ad-hoc leads and lags of consumption. Porter and Wheatley (1999) find that time-aggregation can lead one to estimate a habit where there is none; one could see if time-aggregation leads one to a biased habit estimate given that the true model *does* have habits. We could calculate sampling distributions as well as population values of statistics. Such exercises are of limited value however. First, a consumption-based model is true in our artificial data, by construction. As more complex models approach the assumed model, they are guaranteed to do better. But how good various other consumption-based models are as approximations to *our* model is a question of limited interest; we want to know how good they are as approximations to the truth.

4. Conclusion

We generate artificial time series from a consumption-based model, but the CAPM is a much better approximate asset pricing model than is the canonical power-utility consumption-based model, and multifactor extensions of the CAPM that use price information are better still. We conclude that this finding in real data should be interpreted as evidence against specific functional forms and parameterizations rather than as evidence against consumption-based models in general. Since conditioning information is at the heart of the story, we also conclude more generally that asset pricing models which take account of time-varying conditioning information are likely to perform better than models that do not do so.

We left out measurement error in consumption data, and our model only recognizes a single shock, so that all series are perfectly correlated at high frequency. Generalizing both limitations will provide an even more realistic comparison, and will further degrade the relative performance of consumption-based models in our model economy. On the other hand, we regard it as an interesting success that so much of the relative performance of portfolio-based models can be captured by the effect of conditioning information alone.

Of course, this analysis does not establish that our specification of habit persis-

tence explains the actual cross-section of expected returns based on measured consumption data. Such a test has not yet been performed. There is some hope: Campbell and Cochrane (1999) show that historical consumption data, when fed into the calibrated model, produce stock market swings that are similar in many ways to the actual history of the stock market. However, one must take seriously measurement error and specification error in consumption data before estimating and testing any consumption-based model and especially in comparing it to portfolio-based models.

Furthermore, pricing error comparisons between consumption-based and portfolio-based asset pricing models are fundamentally not that revealing. Returns are far better measured than consumption data, so even if we knew the true utility function, a return-based model (using the mimicking portfolio for marginal utility) would produce smaller pricing errors than the underlying consumption-based model. Such a return-based model would continue to be the best specification for non-structural questions including risk adjustment, anomaly exploration and cost of capital calculations, and for most practitioners, especially at high frequency. Ad-hoc portfolio factors can more closely approximate the ex-post mean-variance efficient portfolio and thus will seem to do even better in any statistical horse race. Consumption-based models will always be used to understand the deeper economic forces that determine the prices of risk in portfolio-based models, to help sort out which ones really work and which were just lucky in particular samples, and to answer structural questions which change the distribution of risks or risk aversion.

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.
- Breeden, Douglas T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics* 7, 265-296.
- Breeden, Douglas T.; Michael R. Gibbons, and Robert H. Litzenberger, 1989, "Empirical Tests of the Consumption-Oriented CAPM," *Journal of Finance* 44, 231-62.
- Campbell, John Y. and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107, 205-251.
- Chen, Nai-Fu, Richard Roll and Stephen A. Ross, 1986, "Economic Forces and the Stock Market," *Journal of Business* 59, 383-403.
- 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., 1999, *Asset Pricing* Manuscript, University of Chicago. Available at <http://www-gsb.uchicago.edu/fac/john.cochrane>.
- Epstein, Larry G. and Stanley E. Zin, 1991, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy* 99, 263-86.
- 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* 47 426-465.
- Ferson, Wayne E. and George M. Constantinides, 1991, "Habit Persistence and Durability in Aggregate Consumption: Empirical Tests," *Journal of Financial Economics* 29, 199-240.
- Grossman, Sanford and Robert J. Shiller, 1981, "The Determinants of the Variability of Stock Market Prices," *American Economic Review* 71, 222-227.
- Hansen, Lars Peter and Kenneth J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica* 50, 1269-1288.

- Hansen, Lars Peter and Kenneth J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy* 91, 249-268.
- Hansen, Lars Peter, and Ravi Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance* 52, 557-590.
- Jagannathan, Ravi and Z. Wang, 1996, "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance* 51, 3-53.
- Kandel, Shmuel, and Robert F. Stambaugh, 1995, "Portfolio Inefficiency and the Cross-Section of Expected Returns," *Journal of Finance* 50, 157-184.
- Lucas, Robert E., Jr. 1978, "Asset Prices in an Exchange Economy," *Econometrica* 46, 1429-45.
- Roll, Richard, and Stephen A. Ross, 1994, "On the Cross-Sectional Relation between Expected Returns and Betas," *Journal of Finance* 49, 101-121.
- Mankiw, N. Gregory and Matthew D. Shapiro, 1986, "Risk and Return: Consumption Versus Market Beta," *Review of Economics and Statistics* 68, 452-259.
- Shanken, Jay, 1987, "Multivariate Proxies and Asset Pricing Relations," *Journal of Financial Economics* 18, 91-110.
- Wheatley, Simon, 1988, "Some Tests of International Equity Integration," *Journal of Financial Economics* 21, 177-212.

</ref_section>