Explaining the Poor Performance of Consumption-based Asset Pricing Models

JOHN Y. CAMPBELL and JOHN H. COCHRANE*

ABSTRACT

We show that the external habit-formation model economy of Campbell and Cochrane (1999) can explain why the Capital Asset Pricing Model (CAPM) and its extensions are better approximate asset pricing models than is the standard consumption-based model. The model economy 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. Therefore, though the consumption-based model and CAPM are both perfect *conditional* asset pricing models, the portfolio-based models are better approximate *unconditional* asset pricing models.

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) show 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 prove disappointing empirically. Hansen and Singleton (1982, 1983) formulate a canonical consumption-based model in which a representative investor has time-separable power utility of consumption. They reject the model on U.S. data, finding that it can not simultaneously explain the time-variation of interest rates and the cross-sectional pattern of average returns on stocks and bonds. Wheatley (1988) rejects 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 performs no better, and in many respects worse, than even the simple static Capital Asset Pricing Model (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. Shanken (1990) Ferson and Schadt (1996), 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. Twenty-five 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 delink 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. The answer is likely to be

deeper than measurement errors in available consumption data sets. 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.

Because 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 (and ignoring the small 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 that 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 is 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.

Section I quickly reviews the Campbell-Cochrane (1999) model and parameter choices. Section II reviews the approximate models and the procedure for calculating their pricing errors. Section III presents the results, and Section IV concludes.

I. 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.

A. Setup

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

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

Identical agents maximize the utility function

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

Here C_t is consumption, X_t is the level of habit, and δ is the subjective or time discount factor. Lowercase letters denote logarithms of uppercase letters, $c_t = \ln C_t$, and so forth. g denotes the mean consumption growth rate.

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

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

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

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

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

where

$$\bar{S} \equiv \sigma \sqrt{\frac{\gamma}{1 - \phi}}, \tag{5}$$

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

$$s_{\max} \equiv \bar{s} + \tfrac{1}{2}(1 - \bar{S}^2). \tag{6}$$

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 risk-free 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 equation (3) like an increase in risk, which increases precautionary savings, thus 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 nonnegatively 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.

B. 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

$$u_c(C_t, X_t) = (C_t - X_t)^{-\gamma} = S_t^{-\gamma} C_t^{-\gamma}. \tag{7}$$

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

$$M_{t+1} \equiv \delta \, \frac{u_c(C_{t+1}, X_{t+1})}{u_c(C_t, X_t)} = \delta \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma}. \tag{8}$$

The log marginal rate of substitution is

$$m_{t+1} = \ln(\delta) - \gamma(\Delta s_{t+1} + \Delta c_{t+1})$$
$$= \ln(\delta) - \gamma g - \gamma(\phi - 1)(s_t - \bar{s}) - \gamma(1 + \lambda(s_t))v_{t+1}. \tag{9}$$

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 risk-free interest rate is the reciprocal of the conditionally expected stochastic discount factor

$$R_t^f = 1/E_t(M_{t+1}). \tag{10}$$

Using equation (9) and taking the expectation of the lognormal random variable M, the log risk-free rate is

$$r_t^f = -\ln(\delta) + \gamma g - \tfrac{1}{2}\gamma(1-\phi).$$ (11)

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

$$\frac{P_t}{C_t}(s_t) = E_t\left[M_{t+1} \frac{C_{t+1}}{C_t} \left(1 + \frac{P_{t+1}}{C_{t+1}}(s_{t+1}) \right) \right].$$ (12)

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

$$\Delta d_{t+1} = g + w_{t+1}; \qquad w_{t+1} \sim \text{i.i.d. } \mathcal{N}(0,\sigma_w^2), \quad \text{corr}(w_t,v_t) = \rho.$$ (13)

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.

C. *Choosing Parameters*

We use the same parameter values as in Campbell and Cochrane (1999), calibrated to postwar (1947 to 1995) annual NIPA nondurable and services per capita consumption together with data from the CRSP value-weighted NYSE stock portfolio. Table I 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 risk-free 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, $\bar{S}$, and $S_{\max}$ follow from these choices via equations (11), (5), and (6).

Parameter Choices

Starred (*) entries are annualized values, for example $12g$, $\sqrt{12}\sigma$, $12r^f$, ϕ^{12}, δ^{12}, $\sqrt{12}\sigma_w$ because the model is simulated at a monthly frequency. δ, $\bar{S}$, and $S_{\max}$ are calculated using equations (11), (5), and (6) respectively.

Panel A: Assumed Parameters	
Mean consumption growth* g, %	1.89
Standard deviation of consumption growth* σ, %	1.50
Log risk-free 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
Panel B: 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

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.

II. 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?

A. False Models

We specify the alternative asset pricing models in terms of their stochastic discount factors, which we denote by 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,

$$Y_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\eta}. \tag{14}$$

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,

$$Y_{t+1} = a + b\left(\frac{C_{t+1}}{C_t}\right). \tag{15}$$

A discount factor that is a linear function of a set of variables implies that expected returns are linear in betas, that is, multiple regression coefficients of returns on those variables. Thus equation (15) corresponds to tests that compare expected returns with consumption betas, without imposing the nonlinear structure of equation (14), as in Mankiw and Shapiro (1986) and Breeden et al. (1989).

(3) The traditional static CAPM,

$$Y_{t+1} = a + bR_{t+1}^w, \tag{16}$$

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

(4) A multifactor or conditional CAPM,

$$Y_{t+1} = [a_0 + a_1(d_t - p_t)] + [b_1 + b_2(d_t - p_t)]R_{t+1}^w \tag{17}$$

$$= a_0 + a_1(d_t - p_t) + b_1 R_{t+1}^w + b_2[(d_t - p_t)R_{t+1}^w]. \tag{18}$$

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

$$y_{t+1} = \ln(\beta) - \eta(c_{t+1} - c_t). \tag{19}$$

The true model is, from equation (9),

$$m_{t+1} = \ln(\delta) + [\gamma g\lambda(s_t) - \gamma(\phi - 1)(s_t - \bar{s})] - \gamma(1 + \lambda(s_t))(c_{t+1} - c_t). \tag{20}$$

The true model has the same form as the canonical consumption-based model (19), except that it makes the intercept and slope on consumption growth functions of the slow-moving state variable s_t. The multifactor model (17) also has this form, but it is driven by the market return rather than by consumption growth.

B. Pricing Errors

We want to know how big the pricing errors of the false models are 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 ratio, 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. We show that

$$\max_{\{\xi\}} \frac{|E^Y(\xi) - E(\xi)|}{\sigma(\xi)} = \frac{\sigma(M - Y)}{E(M)}. \tag{21}$$

The left-hand side of equation (21) 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 risk-free rate and the market return. When the discount factor Y is a linear function of factors, $Y = \mathbf{b'f}$, 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,

$$\min_{\{\mathbf{b}\}} E[(M - \mathbf{b'f})^2] \Rightarrow E(M\mathbf{f}) = E(\mathbf{f}(\mathbf{f'b})) = E(Y\mathbf{f}). \tag{22}$$

Minimizing the maximum pricing error is also equivalent to an OLS regression of the true discount factor M on the factors $\mathbf{f}$. From equation (22),

$$\mathbf{b} = E(\mathbf{ff'})^{-1}E(M\mathbf{f}). \tag{23}$$

Both procedures imply that the true and approximate discount factor agree on the risk-free rate, so $E(M) = E(Y)$, which is required to derive our simple version of Hansen and Jagannathan's distance measure, equation (21).

C. 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

$$E(\xi) = \frac{P}{E(M)} - \frac{\mathrm{cov}(M,\xi)}{E(M)}. \tag{24}$$

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

$$E^Y(\xi) \equiv \frac{P}{E(Y)} - \frac{\mathrm{cov}(Y,\xi)}{E(Y)}. \tag{25}$$

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

$$|E^Y(\xi) - E(\xi)| = \left| \frac{\mathrm{cov}\, n\,(M - Y, \xi)}{E(M)} \right| \leq \frac{\sigma(M - Y)\sigma(\xi)}{E(M)}. \tag{26}$$

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 the Hansen-Jagannathan measure 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 the test portfolio choice.

The Hansen-Jagannathan procedure eliminates this dependence by evaluating the pricing error of the worst possible portfolio, the one that generates the largest possible pricing error. The search for the worst-priced payoff extends over all possible contingent claims, including all dynamic strategies. The worst-priced payoff $\xi^* = M - Y$ is typically 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 nonlinear function of consumption growth, that fact may really not tell us much about which models price stock portfolios well.

Even within the limited set of assets that have been examined, results seem to be sensitive to the asset choice; Breeden et al. (1989) use industry and bond portfolios, and find better results for the canonical consumption-based model than do 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 (21) generalizes to

$$\max_{\{\xi \in \mathbf{X}\}} \frac{|E^Y(\xi) - E(\xi)|}{\sigma(\xi)} = \min_{\{M:\, P=E(MX)\,\forall X \in \mathbf{X}\}} \frac{\sigma(M - Y)}{E(M)}, \tag{27}$$

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 given by equation (13) and pricing them via the true model. One could then evaluate pricing errors of false models via equation (27). 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 equation (13). Therefore, we do not limit the space of portfolios over which to search for large pricing errors.

III. Results

Table II gives the maximum pricing errors in artificial data from the Campbell-Cochrane (1999) model. These pricing errors are estimates of 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 percent average return at a 20 percent 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 percent 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.

Row (c) presents the canonical consumption-based model with power utility. In this row, we estimate the parameters of the model to minimize the maximum pricing error. This is as well as the power-utility model can do,

Table II

Maximum Pricing Errors from Approximate Asset Pricing Models

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 percent 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 risk-free rate. In rows (a), (b), and (c) we minimize the maximum pricing error, which is equivalent to correctly pricing the factors in rows (a) and (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.

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

and it fares a good deal worse than the CAPM. The maximum pricing error is about 30 percent larger than that of the CAPM. We estimate a large coefficient of risk aversion ($\eta = 29$), as one does in real data (Hansen and Singleton (1982)).

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 one almost always picks 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 percent average return for a 20 percent 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 is not 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 II. Recall that the marginal rate of substitution is given by

$$M_{t+1} = \delta \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma}. \tag{28}$$

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 equation (9).

Panel B of Table II 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 to the unconditional pricing of shorter-horizon returns, 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 et al. (1989) and Mankiw and Zeldes (1986). 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. This dividend model is not perfect; its most glaring fault is that dividends and consumption are not cointegrated. However, a more realistic dividend model would require another state variable, and we can at least check the results' sensitivity with this simple model. 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, because *a* consumption-based model is true in our artificial data, by construction. As more complex models approach the assumed model, they will naturally 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.

IV. Conclusion

We generate artificial time series from *a* consumption-based model. In our artificial data, the CAPM, using the wealth portfolio return or the return on a claim to dividends poorly correlated with consumption, is a much better approximate asset pricing model than is the canonical power-utility consumption-based model. 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. Because

conditioning information is at the heart of the story, we also conclude that asset pricing models that take account of time-varying conditioning information are likely to perform better than models that do not do so.

We leave out measurement error in consumption data, and our model only recognizes a single shock, so all series are perfectly correlated at high frequency. Generalizing both limitations provides a more realistic comparison, and should 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 persistence explains the actual cross section of expected returns based on measured consumption data. Empirical work on this issue is only just beginning (Lettau and Ludvigson (1999)). 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 nonstructural 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. Therefore, it is likely that consumption-based models will always be best 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 analyze structural changes in 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, 231–262.

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, 263–286.

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.

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 Rudi W. Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425–461.

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

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.

Jagannathan, Ravi, and Zhenyu 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.

Lettau, Martin, and Sydney Ludvigson, 1999, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, Working paper, Federal Reserve Bank of New York.

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

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

Porter, Robert, and Simon Wheatley, 1999, Time non-separable preferences or artifact of temporal aggregation?, Working paper, University of Florida.

Roll, Richard, and Stephen A. Ross, 1994, On the cross-sectional relation between expected returns and betas, *Journal of Finance* 49, 101–121.

Shanken, Jay, 1987, Multivariate proxies and asset pricing relations, *Journal of Financial Economics* 18, 91–110.

Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99–120.

Wheatley, Simon, 1988, Some tests of international equity integration, *Journal of Financial Economics* 21, 177–212.