Macroeconomic Dynamics, **1**, 1997, 333–354. Printed in the United States of America.

GROWTH-OPTIMAL PORTFOLIO RESTRICTIONS ON ASSET PRICING MODELS

RAVI BANSAL
Fuqua School of Business, Duke University

BRUCE N. LEHMANN
Graduate School of International Relations and Pacific Studies, University of California at San Diego

We show that absence of arbitrage in frictionless markets implies a lower bound on the average of the logarithm of the reciprocal of the stochastic discount factor implicit in asset pricing models. The greatest lower bound for a given asset menu is the average continuously compounded return on its growth-optimal portfolio. We use this bound to evaluate the plausibility of various parametric asset pricing models to characterize financial market puzzles such as the equity premium puzzle and the risk-free rate puzzle. We show that the insights offered by the growth-optimal bounds differ substantially from those obtained by other nonparametric bounds.

Keywords: Asset Pricing Models, Growth-Optimal Portfolio, Lower Bound, Nonparametric Bounds

1. INTRODUCTION

One of the pillars of asset pricing theory is the premise that asset prices permit no arbitrage opportunities. In frictionless markets, this assumption implies the existence of strictly positive state prices that price all traded assets [see Ross (1977, 1978), Harrison and Kreps (1979), Chamberlain and Rothschild (1983), and Hansen and Richard (1987)]. State prices can be transformed into stochastic discount factors, random variables whose realizations are the state prices per unit probability implicit in asset returns, and any asset pricing model is characterized by *its* stochastic discount factor. The linear asset pricing models, such as the Capital Asset Pricing Model (CAPM), the Arbitrage Pricing Theory (APT), and the aggregate consumption-based equilibrium models, that dominate asset pricing theory identify particular stochastic discount factors that price traded assets.

We thank conference participants at the Conference on Computation and Estimation in Finance and Economics held at Washington University on September 15–16, 1995, and, in particular, the conference organizers, William Barnett, Lars P. Hansen, Andrew Lo, and George Tauchen. We owe a special debt to our discussant, Lars Hansen, whose many excellent questions caused us to seek some hard answers, several of which materially changed the paper. Address correspondence to: Ravi Bansal, Fuqua School of Business, Duke University, Durham, NC 27708, USA; e-mail: rb7@mail.duke.edu.

© 1997 Cambridge University Press 1365-1005/97 $9.00 + .10

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

Can asset market data be used to learn about economically interesting attributes of the *set* of stochastic discount factors? Hansen and Jagannathan (1991) provide one such class of restrictions: lower bounds on the variance of the set of stochastic discount factors identified in given asset menus. Snow (1991) shows how to use their approach to bound higher moments as well. Exploring any such necessary implications helps to diagnose parametric asset pricing models and to evaluate their plausibility, to aid in the development of new asset pricing models, and to characterize aspects of asset market data that are puzzling.

We provide a complementary nonparametric bound: a lower bound on the mean of the logarithm of the reciprocal of the set of strictly positive stochastic discount factors. For a given asset menu, the sharpest lower bound is provided by the expected return of the *growth-optimal portfolio*, the portfolio that would be selected by an optimizing investor with log utility restricted to these assets, that maximizes the expected continuously compounded (i.e., log) return.[1] As is the case with the Hansen–Jagannathan (1991) variance bound, the growth-optimal bound is a *necessary* implication of the absence of arbitrage that any valid stochastic discount factor must satisfy.

There are interesting economic intuitions underlying the growth-optimal bound. It requires the mean of the log of the reciprocal of *any* strictly positive stochastic discount factor to equal or exceed the highest continuously compounded growth rate of wealth attainable in asset markets. Accordingly, the reciprocal of marginal utility must grow at least as fast as wealth in equilibrium asset pricing models which, in a variety of parametric representative-agent models, requires the average continuously compounded growth rate of consumption, appropriately scaled for risk aversion, to exceed average compound returns on invested wealth if investors are saving optimally.

This focus on growth dramatically simplifies the analysis of one class of models often invoked to account for the *risk-free rate* and *equity premium* puzzles—habit formation models. Habit formation introduces additional stationary components in intertemporal marginal rates of substitution in these models, the volatility of which facilitate satisfaction of the Hansen–Jagannathan bounds. However, these temporary components do not affect the continuous growth rate of the inverse of marginal utility in these models and hence they do not affect the log bound. That is, the growth-optimal bound in the usual implementations of these models depends only on the growth rate of consumption, time preference, and the relative risk aversion of the investor, and *not* on the parameters that determine the degree of nonseparability over time. We use this special feature of the growth-optimal bound to characterize the plausibility of these models as explanations of the various asset market puzzles.

The plan of the paper is as follows: The next section lays out three definitions of arbitrage opportunities that lead to distinct restrictions on implicit state prices and stochastic discount factors and discusses the Hansen–Jagannathan and growth-optimal bounds. The penultimate section evaluates several economic models from the perspective of growth-optimal bounds and summarizes some of the

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

directions for future research from this growth-optimal perspective. The last section provides concluding remarks.

2. BOUNDS ON BEHAVIOR OF STOCHASTIC DISCOUNT FACTORS

2.1. State Prices and Stochastic Discount Factors Under Different Notions of Arbitrage

This section discusses alternative definitions of arbitrage opportunities and their implications for the behavior of stochastic discount factors in a world with two dates, today and tomorrow. There is a collection of marketed payoffs $\mathcal{X}$ and each random payoff of $\tilde{x} \in \mathcal{X}$ units of account tomorrow trades for $\pi(\tilde{x})$ units of account today. The riskless asset, if one exists, costs $\pi(1)$ units of account today and pays one unit of account tomorrow in each state. Markets are frictionless—there are no taxes, transaction costs, short-sale constraints, or other impediments to free trade. Investors prefer more to less and, hence, would vigorously exploit any perceived arbitrage opportunities. Investors agree on the possible states of nature, a necessary condition for equilibrium in the absence of arbitrage.

What differs across characterizations of stochastic discount factors is the nature of the arbitrage opportunities ruled out in the associated pricing operator. In particular, consider the following three gradations of the definition of the absence of arbitrage opportunities:

Law of one price: All payoffs in $\mathcal{X}$ that pay the same amount in each state of nature have the same price today. That is, $\tilde{x} = \tilde{y} \Rightarrow \pi(\tilde{x}) = \pi(\tilde{y})$.

Weak no-arbitrage: All payoffs in $\mathcal{X}$ that pay nonnegative amounts in all states of nature have weakly positive prices and those that pay nothing in all states of nature have zero prices. That is, $\tilde{x} \geq 0 \Rightarrow \pi(\tilde{x}) \geq 0$ and $\tilde{x} = 0 \Rightarrow \pi(\tilde{x}) = 0$.

Strong no-arbitrage: All payoffs in $\mathcal{X}$ that pay a strictly positive amount in some state of nature and nonnegative amounts in all states of nature have a strictly positive price. That is, $\tilde{x} \geq 0 \Rightarrow \pi(\tilde{x}) > 0$.

The absence of arbitrage restricts the prices in this asset market. If market prices are arbitrage-free in any of the senses enumerated above, there exists a set of state prices, not necessarily positive or unique, that support the linear pricing operator $\pi(\tilde{x})$. Alternatively, the pricing operator can be represented in terms of stochastic discount factors, random variables $\tilde{q}$ whose realizations are given by state prices per unit probability, such that $\pi(\tilde{x}) = E[\tilde{x}\tilde{q} \mid \mathcal{F}]$ for some information set $\mathcal{F}$.[2]

These three definitions place distinct restrictions on the signs of these stochastic discount factors. The *law of one price* places no restrictions on their signs whereas the *weak no-arbitrage* condition implies that the implicit stochastic discount factors are nonnegative but permits them to be zero. The more general *strong no-arbitrage* condition is both a necessary and sufficient condition for the existence of strictly positive state prices.[3]

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

2.2. Variance Bounds

Accordingly, let Q, $Q^+ \subset Q$, and $Q^{++} \subset Q^+ \subset Q$ denote the set of all, nonnegative, and strictly positive stochastic discount factors, respectively. Hansen and Jagannathan (1991, 1994) provide a careful and detailed discussion of the relations among these sets with special reference to their boundaries. In particular, one of the main contributions of Hansen and Jagannathan (1991, 1994) is the construction of lower bounds on the variance of the stochastic discount factors in Q, Q^+, and Q^{++}. All stochastic discount factors satisfy the Euler equation $\pi(\tilde{x}) = E[\tilde{x}\tilde{q}]$ and, hence, any stochastic discount factor that is also a payoff in $\mathcal{X}$ is a mean-variance frontier portfolio when suitably scaled. This permits Hansen and Jagannathan to construct minimum variance stochastic discount factors that are contained in Q and Q^+.

For any vector of payoffs $\tilde{x} \in \mathcal{X}$ that includes a unit payoff, the minimum variance stochastic discount factor in Q is the unique portfolio p^* with weights $E[\tilde{x}\tilde{x}']^{-1}\pi(\tilde{x})$. Given the minimum variance property of p^* any hypothetical stochastic discount factor $\tilde{q} \in Q$, must have a variance at least as large as that of p^* because $E(\tilde{q}) = E(p^*)$ with a unit payoff. Similarly, the nonnegative minimum variance discount factor is the payoff p^+ constructed by finding a return $\tilde{R}_p^+ \in \mathcal{X}$ with the smallest truncated second moment $E[\max(0, \tilde{R}_p^+)^2]$ and then scaling $\max[0, \tilde{R}_p^+]$ by $E[\max(0, \tilde{R}_p^+)^2]$. The volatility of this nonnegative minimum variance discount factor $p^+ \in Q^+$ provides a lower bound on the volatility of any hypothesized proxy $\tilde{q}$ in Q^+ because $E(\tilde{q}) = E(p^+)$ in the presence of a unit payoff.[4] Hansen and Jagannathan (1991) also establish that the variance of p^+ generally will exceed that of p^*, the difference in the minimum variance bound reflecting the incremental contribution of imposing the restriction of *weak no-arbitrage*. Because Q^+ is the mean-squared closure of Q^{++}, the variance of p^+ is a lower bound on the variance of all stochastic discount factors in Q^{++} as well. The Hansen–Jagannathan bounds exhaust all necessary implications for the *variance* of a discount factor hypothesized to be in Q, Q^+, and Q^{++}.

Cochrane (1992) and Stutzer (1995) provide alternative nonparametric bounds on stochastic discount factors. Cochrane (1991) focuses on nonparametric bounds that relate to the variance of price dividend ratios and notes that there are associated restrictions on the mean log discount rate {i.e., $-E[\log(\tilde{q})]$} must be at least as large as the average of any continuous dividend growth rate and the average of any continuous unit cost return to avoid an infinite price dividend ratio.[5] Stutzer (1995) provides a nonparametric bound on the Kullback–Leibler Information Criterion associated with the discount factor. Recently, Bernardo and Ledoit (1996) extended growth-optimal and exponential utility bounds to the entire constant relative risk aversion class. Each of these nonparametric bounds provides insights regarding different aspects (moments) of proxy discount factors.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

2.3. Growth-Optimal Bounds

Let $\tilde{\boldsymbol{R}} \in \mathcal{X}$ denote the vector of gross returns (i.e., unit cost payoffs) and let w be an arbitrary set of portfolio weights that sum to one (i.e., cost one unit of account to scale the problem). For any $\tilde{q} \in Q^{++}$, the relevant pricing restriction is

$$E[\tilde{q}\ \ w'\tilde{R}] = 1, \qquad w'\iota = 1 \tag{1}$$

or, equivalently, that

$$\tilde{q}\ \ w'\tilde{R} = 1 + \tilde{u}_{wq}; \qquad E[\tilde{u}_{wq}] = 0. \tag{2}$$

Now, confine attention to portfolios with strictly positive payoffs with probability 1. Taking the expected value of the natural logarithm of $\tilde{q}\ \ w'\tilde{R}$ in this case yields

$$E[\ln(\tilde{q})] + E[\ln(w'\tilde{R})] = E[\ln(1 + \tilde{u}_{wq})] \leq \ln E[1 + \tilde{u}_w q] = 0, \tag{3}$$

where the weak inequality exploits Jensen's inequality. As noted above, Cochrane (1992) also recognizes the inequality in (3).

Searching for the tightest log bound reveals a special role for the *growth-optimal portfolio*. Maximization of the right-hand side of (3) yields

$$-E[\ln(\tilde{q})] \geq \max_{\{w_g\}} E[\ln(w_g'\tilde{R})] \quad \text{s.t. } w_g'\iota = 1, \qquad \tilde{q} \in Q^{++}, \tag{4}$$

so that the growth-optimal portfolio provides the sharpest lower bound by construction.[6] Hence, the *growth-optimal bound* is

$$-E[\ln(\tilde{q})] \geq E[\ln(\tilde{R}_g)] \tag{5}$$

because the return $\tilde{R}_g \equiv w_g'\tilde{R}$ is strictly positive because the log investor never risks ruin for any state of the world—that is, $1/\tilde{R}_g > 0$.[7]

The nonparametric bound in (5) makes economic sense as well when $\tilde{q}$ is interpreted as the intertemporal marginal rate of substitution of some optimizing investor. In this case, it captures the idea that the inverse of marginal utility must grow at least as fast as the maximal attainable growth rate of wealth for $\tilde{q}$ to be the stochastic discount factor of such an investor. Hence, the underlying economic intuition differs substantially from that of the Hansen–Jagannathan bound.

In many cases, it is also convenient to also have an excess return version of the bound (5). The construction of such a bound is straightforward when $\mathcal{X}$ includes a unit payoff or when it is augmented with one. Such a bound is obtained by adding $\ln E[\tilde{q}]$ and noting that this quantity equals the mean price of a riskless bond $\ln(P_f)$. Hence,

$$-E[\ln \tilde{q}] + \ln E[\tilde{q}] \geq E[\ln(\tilde{R}_g)] + \ln(P_f). \tag{6}$$

As is readily apparent, the role of the risk-free bond price is to scale the stochastic discount factor to sharpen the focus on risk premiums. Given a proxy discount

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

factor, the left-hand side of (6) is the excess-return on the proxy with the highest risk premium and the right-hand side is the highest risk premium observed in the asset market ($\mathcal{X}$). To satisfy (6), the risk premium on the proxy must exceed the highest risk premium observed in asset returns. The excess-return bound also eliminates the effect of the rate of time preference on the bound, which is useful in some applications.

The excess-return bound also clarifies the circumstance in which the Hansen–Jagannathan and growth-optimal bounds are similar: the case of lognormality. Accordingly, suppose that $\tilde{R}_g^{-1}, \tilde{q} \in Q^{++}$ are each lognormally distributed. In this case, the excess-return version of the growth-optimal bound becomes a minimum variance bound:

$$\text{Var}[\ln(\tilde{q})] \geq \text{Var}[-\ln(\tilde{R}_g)] \tag{7}$$

because $E[\tilde{q}] = E[\tilde{R}_g^{-1}]$ and $E[\ln(x)] = \ln[E(x)] - \frac{1}{2}\text{Var}[\ln(x)]$. Consequently, the growth-optimal bounds under this distributional assumption provides almost the same information as the minimum variance bounds, particularly in the continuous-time diffusion limit.

Simple manipulation of (6) reveals the role played by the moments of $\tilde{q}/E(\tilde{q})$ in it. Its left-hand side is

$$-E\left\{\ln\left[1 + \frac{\tilde{q}}{E(\tilde{q})} - 1\right]\right\} \quad \text{and let} \quad \frac{\tilde{q}}{E(\tilde{q})} - 1 \equiv \eta,$$

where η is mean zero by construction. Expanding $\ln[1 + \eta]$, and taking its expectation implies

$$-E\left\{\sum_{i=1}^{i=\infty} \frac{-1^{i-1}}{i}[\eta]^i\right\} \geq E[\ln(\tilde{R}_g)] - \ln(R_f)], \qquad E[\tilde{q}] = E\left[\frac{1}{\tilde{R}_g}\right]. \tag{8}$$

Equation (8) shows that the excess-return bound is determined by all of the central moments of $\tilde{q}/E(\tilde{q})$.[8] In particular, large even moments and small odd moments facilitate the satisfaction of the growth-optimal bound. Hence, negatively skewed stochastic discount factors with excess kurtosis attain the bound more readily than symmetrically and approximately normally distributed ones with the same variance. Similarly, comparison of (7) for the lognormal case with the more general expression (8) serves as a reminder that the growth-optimal and Hansen–Jagannathan bounds are likely to differ markedly only to the extent that lognormality is a bad approximation for returns and stochastic discount factors, a prospect that we think likely given the prevalence of comparatively infrequent but severe crashes in financial markets coupled with more common but more modest large positive returns.

There are two additional features of growth-optimal portfolios that we have not exploited above: their myopia and numeraire invariance. The growth-optimal strategy is myopic in the sense that an investor who maximizes the discounted

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

expected utility of the log of consumption over future dates will find it optimal to solve the single-period program each period. This feature simplifies the link between single-period and multiperiod stochastic discount factors, facilitating the derivation of multiperiod growth-optimal bounds that can focus more sharply on low-frequency attributes of the set of stochastic discount factors. By contrast, the multiperiod analogues of p^* and p^+ generally are not the associated products of the single-period ones. Numeraire invariance follows because the logarithm of a product is the product of the logarithms and a change in numeraire is just a change in scale, making the growth-optimal portfolio weights identical irrespective of whether the maximum problem is stated in real or nominal terms. This feature is convenient empirically because we can use nominal returns, which are well measured, as opposed to real returns, which suffer from errors in the measurement of inflation rates. By contrast, the portfolio weights implicit in p^* and p^+ generally depend on whether they are computed in real or nominal terms.

3. GROWTH-OPTIMAL BOUNDS AND PUZZLES

We now illustrate some of the empirical regularities that arise from the growth-optimal bound. Our focus is on documenting and interpreting the shortcomings of commonly investigated asset pricing models from the perspective of the growth-optimal bound. In particular, we show how our bound provides insights into models based on both time-separable and time-nonseparable preferences that are distinct from those derived from the Hansen–Jagannathan bounds.

We also should be clear about the scope of the investigation conducted in this paper. The objective is to provide evidence that constitutes a prima facie case for the usefulness of this approach. A more extensive analysis would explore a wider range of asset menus, conditioning information, additional candidate stochastic discount factors, and the impact of violating the assumption of frictionless markets. Such an investigation also would explore some of the methodological and statistical issues that naturally arise in this setting, particularly those related to the small-sample properties of growth-optimal portfolios that exploit nontrivial conditioning information.

3.1. Estimation and Inference Issues

Let M be the hypothesized stochastic discount factor implied by some asset pricing model. In computing the growth-optimal bounds, we set $E(M) = P_f = 1/R_f$ and, in an abuse of language, we sometimes refer to $R_f = 1/P_f$ as the real riskless rate even though it is actually the inverse of the average price of the riskless bond. For this value of P_f, we estimate the growth-optimal weights by applying the Generalized Method of Moments as in Hansen and Singleton (1982) to the first-order conditions of the growth-optimal problem $E[(\tilde{\boldsymbol{R}} - R_f\iota)/\tilde{R}_g] = \mathbf{0}$. Note that the number of such first-order conditions equals the number of excess returns and, hence, the number of unknown portfolio weights in $\tilde{R}_g$. We estimate $E[\ln \tilde{R}_g]$

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

from the corresponding time-series average of $\ln \tilde{R}_g$, which takes the form $\ln[R_f + \hat{w}_g'(\tilde{\boldsymbol{R}} - R_f\iota)]$, where $\tilde{\boldsymbol{R}}$ is the vector of unit-cost real returns used in estimating $\tilde{R}_g$ and $R_f = 1/E(M)$. Note that the growth-optimal portfolio satisfies $E[1/\tilde{R}_g] = E(M)$ by construction, making the bound

$$-E[\ln(M)] \geq E[\ln(\tilde{R}_g)], \qquad E(M) = E(1/\tilde{R}_g). \tag{9}$$

Statistical inference for the mean of the difference $-\ln M_t - \ln \tilde{R}_{gt}$ is straightforward. The unknown parameters of interest are $\theta_0' = \{E[-\ln M_t - \ln \tilde{R}_{gt}], E(M_t), w_g'\}$, also let the parameter subvector $\theta_{0,s}'$ be $\{E[-\ln M_t - \ln \tilde{R}_{gt}], E(M_t)\}$. When N excess returns are used to estimate the growth-optimal portfolio, w_g' is a vector of length N and, consequently, θ_0' is $2 + N$. Consider the $2 + N$ mean zero-error vector $\boldsymbol{u}_t(\theta_0)' = \{-\ln M_t - \ln \tilde{R}_{gt} - E[-\ln M_t - \ln \tilde{R}_{gt}], M_t - E(M_t), (\{\tilde{R}_t - [1/E(M_t)]\}/\tilde{R}_{gt})'\}$. The number of unknown parameters in θ_0 is exactly equal to the number of errors and hence the system is exactly identified. The unknown-parameter vector can be estimated by using the Generalized Method of Moments of Hansen (1982). For this exactly identified system $\sqrt{T}(\hat{\theta}_0 - \theta_0)$ is asymptotically distributed as $\mathcal{N}[0, (d_0^{-1} S_0 d_0^{-1\prime})]$, where $d_0 = E[\partial \boldsymbol{u}_t(\theta_0)/\partial\theta_0]$ and S_0 is the spectral density of the errors at frequency zero [see Hansen (1982)].

In this case, d_0 and S_0 are $(2 + N) \times (2 + N)$ square matrices, and d_0 has the special structure

$$d_0 = \begin{bmatrix} \begin{pmatrix} -1 & \dfrac{[1 - w_g'\iota]}{E(M_t)} \\ 0 & -1 \end{pmatrix} & \mathbf{0} \\ d_{12} & d_{22} \end{bmatrix}.$$

The upper-right partition $\mathbf{0}$ is a matrix of zeros, reflecting the fact that the derivative of the first two errors with respect to the portfolio weights w_g' are zero in the population. Let d_{11} and S_{11} be the upper-left 2×2 partition of d_0 and S_0, respectively. Exploiting the special structure of d_0, it follows that $\sqrt{T}(\hat{\theta}_{0,s} - \theta_{0,s})$ is asymptotically distributed as $\mathcal{N}[0, (d_{11}^{-1} S_{11} d_{11}^{-1\prime})]$. Consequently, if s_{ij} is the ij element of S_{11}, then the asymptotic variance of estimated mean of $-\ln M_t - \ln R_{gt}$ is

$$s_{11} + \left(\frac{[1 - w_g'\iota]}{E(M_t)}\right)^2 s_{22} + 2\frac{[1 - w_g'\iota]}{E(M_t)} s_{12}.$$

Similarly, s_{22} is the asymptotic variance of the estimated mean of M_t. The uncertainty with which the portfolio weights are estimated does not alter the asymptotic distribution of the growth-optimal bound—that is, the estimated mean of $-\ln M_t - \ln R_{gt}$. This implication of our analysis is analogous to a result of Hansen et al. (1995) who discuss it in the context of minimum-variance bounds. Note that the asymptotic variance of the estimated mean of $-\ln M_t - \ln R_{gt}$ is simply s_{11} when $E(M_t)$ is postulated a priori and not estimated. The term $[1 - w_g'\iota]$ is the portfolio weight assigned to the payoff $1/E(M_t)$, which can be quite large when

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

$P_f = E(M_t)$ approaches its arbitrage bound and is unbounded outside the arbitrage bound for P_f. Consequently, the incremental contribution of this term to the asymptotic variance of the estimated mean of $-\ln M_t - \ln R_{gt}$ can be substantial even if the variance of the estimated mean of M_t is small.

3.2. Data and the Bounds

Our monthly data for returns run from January 1959 to July 1991, giving us 390 observations. The two returns used in our empirical exercise are those of the one-month Treasury bill and the CRSP value-weighted index which we converted to ex post real returns using the inflation rate based on the CPI. We augment these assets with a hypothetical one-period real riskless bond at several assumed average real prices to facilitate comparison with Hansen and Jagannathan (1991, 1994) and Cochrane and Hansen (1992).

Table 1 provides summary statistics for the asset returns used in constructing the alternative bounds, and for p^*, p^+, and $1/\tilde{R}_g$. The asset payoffs used in constructing these alternative discount factors are the real value-weighted return and the real T-bill return. The volatility of p^* and p^+ and the average continuous return on $\tilde{R}_g$ take their minimum values when $1/P_f$ is about 1.2%, an unsurprising outcome given that the average real return on the T-bill is about 1.2% per annum. The values of $1/P_f$ that we consider range from -8.4% to 12.00% as shown in Table 2.[9] Outside of this range, the computation of the expected utility of the hypothetical log investor reveals that high and low values of P_f cause violations of the second-order conditions. That is, real prices outside this range lead to empirical arbitrage opportunities as the hypothetical real riskless asset becomes a dominated or dominating asset and the portfolio weights of the growth-optimal portfolio grow without bound. This narrow range sharply restricts the parameters of models of

TABLE 1. Data[a] sample moments

Variable[b]	Mean	Std	Min	Max
c	0.0025	0.0043	0.0180	−0.0134
π	0.0004	0.0030	0.0136	−0.0047
r	0.0009	0.0027	0.0104	−0.0074
vw	0.0041	0.0443	0.1513	−0.2468
p^*	0.9991	0.0957	1.4881	0.6514
p^+	0.9991	0.0957	1.4881	0.6514
$1/\tilde{R}_g$	0.9991	0.0981	1.8199	0.7493

[a] The monthly data are from February 1959 to July 1991.
[b] c is real consumption growth (nondurables plus services); π is inflation; r is the ex-post real return on a one-month T-bill; vw is the real value-weighted index return; p^*, p^+ refer to the estimated minimum-variance stochastic discount factors of Hansen and Jagannathan; and $1/\tilde{R}_g$ is the growth-optimal stochastic discount factor. The various discount factors are estimated using r and vw. The correlation between p^* (or p^+) and $1/\tilde{R}_g$ is 0.973.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

TABLE 2. Estimated growth-optimal bounds

P_f[a]	r[b]	$Std[1/\tilde{R}_g]$[c]	$E[\ln(\tilde{R}_g)]$[c,d]	Standard error[e]
1.0073	−8.7000	14.707	3.9874	0.0480
1.0065	−7.8000	13.021	2.1239	0.0421
1.0058	−6.9000	11.417	1.5157	0.0396
1.0047	−5.7000	9.2909	1.0441	0.0372
1.0037	−4.5000	7.0847	0.7275	0.0357
1.0025	−3.0000	4.2679	0.4384	0.0328
1.0010	−1.2000	1.4225	0.1807	0.0284
1.0000	0.0000	0.4241	0.0543	0.0197
0.9993	0.9000	0.1150	0.0068	0.0063
0.9988	1.5000	0.1622	0.0133	0.0096
0.9977	2.7000	1.0415	0.0999	0.0238
0.9967	3.9000	2.9727	0.2226	0.0272
0.9960	4.8000	4.4262	0.3299	0.0293
0.9955	5.4000	5.3291	0.4102	0.0307
0.9943	6.9000	7.3335	0.6503	0.0342
0.9935	7.8000	8.4336	0.8306	0.0360
0.9927	9.0000	9.8737	1.1387	0.0384
0.9916	10.2000	11.3556	1.5887	0.0411
0.9901	12.0000	13.736	3.5365	0.0498

[a] Chosen mean of the discount factor. The quoted range for P_f also provides the arbitrage bounds for the price of the real riskless payoff.
[b] $r = (1/P_f - 1) * 1200$.
[c] $\tilde{R}_g$ is the estimated growth-optimal portfolio return.
[d] $E[\ln(\tilde{R}_g)]$ is the mean of the continuously compounded growth-optimal return.
[e] Standard error of $E[\ln(\tilde{R}_g)]$ constructed using the Newey–West procedure with two lags.

the stochastic discount factor to those that imply mean prices for the hypothetical real riskless bond that are sufficiently close to unity.

3.3. Time- and State-Separable Preferences

Perhaps the most extensively used class of models in economics is that in which the representative household is assumed to have state- and time-separable preferences of the constant relative risk aversion variety. The stochastic discount factor in this model is

$$M_{t+1} = \beta \lambda_{t+1}^{-\alpha}, \tag{10}$$

where $\lambda_{t+1} = c_{t+1}/c_t$ is the gross growth of aggregate consumption growth, α ($\alpha \geq 0$) is the relative-risk-aversion parameter, and β determines the rate of time preference. Using (6), it follows that the growth-optimal bound for this model is

$$-E[\ln(M_{t+1})] = -\ln(\beta) + \alpha E[\ln(\lambda_{t+1})] \geq E[\ln(\tilde{R}_{gt})], \qquad E(M_t) = E(1/\tilde{R}_{gt}). \tag{11}$$

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

TABLE 3. Time-separable model

α^a	$E[M]^b$	r^c	$E[\ln(\tilde{R}_g)]^d$	$E[\ln(1/M)]$	$Diff^e$	Standard error[f] 1	Standard error[f] 2
1.000	0.9975	2.9976	0.1283	0.0025	−0.1258	0.0182	0.0337
1.500	0.9962	4.4980	0.2923	0.0038	−0.2885	0.0209	0.0538
2.000	0.9950	5.9895	0.4972	0.0050	−0.4922	0.0236	0.0850
2.500	0.9938	7.4771	0.7620	0.0063	−0.7557	0.0262	0.1342
3.000	0.9926	8.9607	1.1270	0.0075	−1.1194	0.0285	0.2229
3.500	0.9913	10.4403	1.7593	0.0088	−1.7505	0.0313	0.4665
4.000	0.9902	11.9160	3.5210	0.0100	−3.5109	0.0380	2.9859
238.0	1.0004	−0.4387	0.0950	0.5970	0.5021	0.0509	10.561
$r - CAPM^g$	0.9969	3.7662	0.2078	0.0041	−0.2037	0.0196	0.3124

[a] $M_{t+1} = [c_{t+1}/c_t]^{-\alpha}$.
[b] Unconditional mean of M.
[c] $r = [1/E(M) - 1] * 1200$ is the inverse of the price of a real riskless bond.
[d] Mean of the continuously compounded growth-optimal return $\tilde{R}_g$.
[e] *Diff* $E[\ln(1/M)] - E[\ln(\tilde{R}_g)]$.
[f] All standard errors are constructed using the Newey–West procedure with two lags: (1) standard error of *Diff* ignoring the sample error in estimating the mean of M, (2) standard error of *Diff* that incorporates the sampling error in estimating the mean of M.
[g] Rubinstein version of the CAPM, where $M = 1/(1 + vw)$.

Table 3 provides information on the moments of this stochastic discount factor and its logarithm when $\beta = 1$ and with α ranging from 1 (log utility) to 4 and $\alpha = 238$ using nondurable plus services consumption from Citibase as the consumption measure. Note that when α is set at 4, the implied price of the real riskless bond becomes very close to the arbitrage bound. In this sample, the average continuously compounded growth rate of real consumption is about 0.0025% per month (3% per annum) and its standard deviation is about 0.0043% per month, making the variance of consumption growth more than two full orders of magnitude smaller than its mean.

For values of α up to 4, the point estimates of $E[\ln(\tilde{R}_{gt})]$ are roughly two full orders of magnitude larger than those of $-E[\ln(M_t)]$. Put differently, the mean values of $-E[\ln(M_t)]$ are below 1% per month for these models whereas the average continuously compounded growth-optimal portfolio return exceeds 12% per month. For values of α up to 3.5, the asymptotic t-ratios (based on the standard errors labeled 2) of the mean difference exceed 3.7 in absolute value, implying rejection of the model using one-sided tests at any conventional significance level. However, the asymptotic standard errors are so large for $\alpha = 4$ that the corresponding asymptotic t-ratio is −1.1758. The point estimate of the bound roughly doubles when α is raised from 3.5 to 4, but the standard error rises by nearly a full order of magnitude. As mentioned earlier, the portfolio weight assigned to the payoff $1/E(M_t)$ increases considerably as the implied price of the riskless bond approaches the arbitrage bound, raising the contribution of the sampling error in the sample mean of the stochastic discount factor.[10] We calculated the

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

standard error ignoring the sampling variation in the mean of M_t (labeled standard error 1); we would sharply reject the model with $\alpha = 4$ were we to use this standard error.

Table 3 also provides information regarding the Rubinstein CAPM, a log utility model in which $M_{t+1} = 1/R_{m,t+1}$ where $R_{m,t+1}$ is the gross return on the value-weighted market portfolio. Treating the CRSP value-weighted index as the market portfolio, the sample mean of $\ln(\tilde{R}_{gt})$ for the Rubinstein CAPM is two orders of magnitude larger than that of $-\ln(M_t)$. However, the standard error on the difference is so large that the asymptotic t-ratio is only -0.66. A comparison of the large-sample standard errors in the last column, which account for the imprecision with which $E[M_t]$ is estimated, with those in the penultimate column, which ignore this estimation error—that is, implicitly assuming that the sample mean of M is at its probability limit—also reveals that sampling error in the mean of M markedly increases the standard error of the bound. In the case of the Rubinstein CAPM, this increase is partly due to the large volatility of the market return, resulting in imprecise estimates of $E(M_t)$. The substantial impact of sampling error in the mean of M_t and the large weight given to the real riskless bond when $\alpha = 4$ indirectly reinforces our a priori belief that the finite-sample properties of growth-optimal portfolio weights are worthy of further investigation.

The time-separable model satisfies the growth-optimal bound for values of α that exceed 230. For example, the point estimate of $E[\ln(\tilde{R}_{gt})]$ is one-sixth that of $-E[\ln(M_t)]$ when $\alpha = 238$, a reversal of the sign in the difference between the two that results in failure to reject this version of the model at any conventional significance level. Note also that accounting for the effect of sampling error in $E[M_t]$ on this difference raises the asymptotic standard error by more than two full orders of magnitude. To a first approximation, the effect of volatility on the mean of this stochastic discount factor dominates that of consumption growth for very high values of α, increasing the real risk-free bond price. Moreover, this increase in α also raises the average continuous return on M_t, which helps to satisfy the growth-optimal bound as well. As is the case in Hansen and Jagannathan (1991) and Cochrane and Hansen (1992), extreme risk aversion facilitates attainment of the growth-optimal bound.[11] However, the interpretation of this result is somewhat different from the perspective of the growth-optimal bound. Extreme risk aversion increases the variance of M considerably, which permits this model to satisfy the minimum variance bounds of Hansen and Jagannathan. By contrast, extreme risk aversion is necessary to increase $-E[\ln(M_{t+1})]$ in the growth-optimal case. Evidently, for this particular model, extreme values of α increase both the variance and the average continuous return on the proxy in this model.

One dimension in which the log and Hansen–Jagannathan bounds deliver different inferences concerns the subjective rate of time preference. Kocherlakota (1990) noted that values of $\beta > 1$ are admissible in equilibrium models when there is positive growth in consumption. Increases in β reduce the implied riskless rate and increase the volatility of the stochastic discount factor in this model

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

(i.e., $\text{var}(M_t) = \beta^2 \text{ var}[\lambda_t^{-\alpha}]$), making it easier to satisfy the Hansen–Jagannathan variance bound. However, increasing β *worsens* the growth-optimal bound in this model because the bounded quantity is the mean of the log of the inverse of the stochastic discount factor; that is, values of β in excess of one make $-\ln(\beta)$ negative. For example, annual discount factors of 1.05 and 1.1 reduce the left-hand side of the bound by 0.0041 and 0.0079, respectively, corrections of the same order of magnitude as the low sample estimates of $\alpha E[\ln(\lambda_t)]$ produced by the smaller risk aversion coefficients reported in Table 3. Salvation for the standard model in the form of negative rates of time preference does not arise when measuring its performance against the growth-optimal bound.

3.4. Time-Nonseparable Preferences

Ryder and Heal (1973), Dunn and Singleton (1986), Sundaresan (1989), Abel (1990), Constantinides (1990), DeTemple and Zapatero (1991), Ferson and Constantinides (1991), Gallant et al. (1990), Heaton (1995), Campbell and Cochrane (1995), and Gali (1994) take the view that time-nonseparable preferences can help resolve asset pricing regularities that are puzzling from the perspective of representative investor models. Hansen and Jagannathan (1991) and Cochrane and Hansen (1992) evaluate some versions of these models with the Hansen–Jagannathan minimum-variance bounds. We show that the growth-optimal bound provides different insights into the plausibility of this class of models.

Consider the representative investor preference specification

$$E\left[\sum_{t=0}^{t=\infty} \beta^t \left(s_t^{1-\alpha}\right) - 1\right] \Big/ 1-\alpha, \qquad s_t = c_t + \sum_{j=0}^{j=J} \theta_j c_{t-j}, \qquad J > 0 \quad \textbf{(12)}$$

in which different models place different restrictions on s_t. Positive values of θ_j imply that consumption provides a service flow whereas negative coefficients capture the notion of habit persistence. A model incorporates durability of consumption at some lags and habits at others when some of the θ_j terms are positive and others are negative.

As is standard, assume that the growth rate of consumption is a strictly stationary process. The pricing condition in this model is

$$E_t[mu_t] = \beta E_t[mu_{t+1} R_{t+1}], \qquad mu_t = s_t^{-\alpha} + \sum_{j=1}^{j=J} \beta^j \theta_j s_{t+j}^{-\alpha}, \qquad \textbf{(13)}$$

where mu_t is the ex post marginal utility of consumption at time t. To preserve strong no-arbitrage, assume that the underlying parameters satisfy $mu_t > 0$ for all states of the economy. The growth-optimal bound is more easily explicated in terms of the transformed variables $mu_t = c_t^{-\alpha} mu_t / c_t^{-\alpha}$ and $h_t \equiv mu_t / c_t^{-\alpha}$. Given the assumed preferences and the process for s, it is easy to show that h_t depends only on past and future growth rates of consumption to some power and, hence,

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

is a strictly stationary process.[12] In terms of the transformed variables, the Euler equation is given by

$$1 = \beta E_t\left[c_{t+1}^{-\alpha} h_{t+1} R_{t+1}\right] / E_t\left[c_t^{-\alpha} h_t\right] = E_t\left[\lambda_{t+1}^{-\alpha} h_{t+1} R_{t+1}\right] / E_t[h_t], \tag{14}$$

and application of the law of iterated expectations yields

$$1 = \beta E_t\left[\lambda_{t+1}^{-\alpha} R_{t+1} E_{t+1}(h_{t+1})\right] / E_t[h_t]. \tag{15}$$

Hence, the discount factor that prices all returns in this economy is

$$M_{t+1} = \beta \lambda_{t+1}^{-\alpha} E_{t+1}(h_{t+1}) / E_t(h_t). \tag{16}$$

Evidently, the one-step-ahead discount factor depends on the growth rate of consumption and the conditional expectation at $t + 1$ and t of the stationary random variable h. Hence, minus the logarithm of the intertemporal marginal rate of substitution in this model is

$$-\ln(M_{t+1}) = -\ln(\beta) + \alpha \ln(\lambda_{t+1}) + \ln\{E_{t+1}(h_{t+1}) / E_t[h_t]\}. \tag{17}$$

Note that h is stationary; further assuming that $E_t[h_t]$ is also stationary implies that the unconditional mean of $\ln\{E_{t+1}(h_{t+1}) / E_t[h_t]\}$ is zero.

Surprisingly, the growth-optimal bound in this model collapses to

$$-E[\ln(M_{t+1})] = -\ln(\beta) + \alpha E[\ln(\lambda_{t+1})] \geq E[\ln(\tilde{R}_{gt})], \tag{18}$$

independent of the parameters that determine the degree of time-nonseparability, although these parameters do affect the average risk-free bond price implied by the model. That is, the growth-optimal bound in internal habit models collapses to that of the standard consumption-based model save for the freedom to use the nonseparability parameters to set the discount bond price essentially at any level.

Note that this statement is only true of the raw return bound (5), not of the excess return bound (6). Evaluation of the latter requires the computation of the average real bond price implied by the model which is a function of all of its parameters. In particular, the nonseparability parameters can serve to enhance the volatility of the log of this stochastic discount factor and, hence, facilitate satisfaction of the excess return bound (6) when the lognormal approximation (7) is a good one.

Abel (1990) and Campbell and Cochrane (1995) discuss external habit models in which the representative investor treats the habit x_t as an exogenous stock that depends on lagged consumption. In the Abel (1990) model, period marginal utility depends on $s_t = c_t / x_t$ while it depends on $s_t = c_t - x_t$ in the Campbell and Cochrane (1995) specification. The period marginal utility for the Campbell and Cochrane (1995) specification is $c_t^{-\alpha}[1 - (x_t / ct)]^{-\alpha} \equiv c_t^{-\alpha} S_t^{-\alpha}$, where S_t is the surplus consumption/habit ratio, making the stochastic discount factor in this model

$$M_{t+1}^{cc} = \beta \lambda_{t+1}^{-\alpha} (S_{t+1} / S_t)^{\alpha}, \tag{19}$$

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

and the corresponding growth-optimal bound is

$$-E\left[\ln M_{t+1}^{cc}\right] = -\ln(\beta) + \alpha E[\ln \lambda_{t+1}] \geq E[\ln(\tilde{R}_{gt})], \quad E\left(M_{t+1}^{cc}\right) = E(1/\tilde{R}_{gt}), \tag{20}$$

because $E[\ln(S_{t+1}/S_t)] = 0$ because S is a stationary process. Once again, the growth-optimal bound for this model is identical to that of the standard consumption-based model. Abel (1990) models the habit as "catching up with the Joneses," in which the period marginal utility is determined by the ratio of current consumption to lagged aggregate consumption, leading to the stochastic discount factor,

$$M_{t+1}^{a} = \beta \lambda_{t+1}^{-\alpha} / \lambda_t^{1-\alpha}, \tag{21}$$

and the corresponding growth-optimal bound

$$-E\left[\ln M_{t+1}^{a}\right] = -\ln(\beta) + E[\ln(\lambda_{t+1})] \geq E[\ln(\tilde{R}_{gt})], \qquad E\left(M_t^{a}\right) = E(1/\tilde{R}_{gt}). \tag{22}$$

The growth-optimal bound in this model is simply that of the additively separable log utility model because the unconditional mean of $\ln[\lambda_{t+1}]$ is the same as that of $\ln[\lambda_t]$.

Hence, the raw-return version of the growth-optimal bound (5) for both of these external habit specifications is identical to that of the corresponding time-separable specification save for differing implications for real riskless bond prices. The more general habit formation specification in the Campbell and Cochrane (1995) model provides the freedom to match the Euler equation associated with the risk-free bond price.[13] The particular external habit structure in the Abel (1990) model implies a set of riskless bond prices associated with different coefficients of relative risk aversion that differ from that imputed from the additively separable log utility consumption CAPM. When the parameters of either model are set to match the real risk-free bond price, the highest continuous return postulated in these models is still determined by the parameters of the corresponding time-separable model—that is, α, β, and the average continuous growth rate of consumption—whereas the risk-free bond price is determined by all parameters of the model, including those that determine the degree to which preferences are not separable over time. This holds for all of the time-nonseparable models that have received considerable attention in the literature. In particular, time-nonseparable preferences with arbitrary lag structures in the model for the habit can explain the various puzzles *only* through their impact on the real risk-free bond price implied by the model when viewed from the perspective of the growth-optimal bound (5).

Given that the left-hand side of the growth-optimal bound is identical across all models for given β and α, we need only consult Table 3, from which the relevant data are displayed in Panel A of Table 4, to determine what the growth-optimal bound based on the return on the T-bill and the value-weighted index has to say about the Campbell and Cochrane (1995) variant of time-nonseparable preference models. The left-hand side of the growth-optimal bound is just the

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

TABLE 4. Habit formation model

α[b]	$E[M]$[c]	r[d]	$E[\ln(\tilde{R}_g)]$[e]	$E[\ln(1/M)]$	*Diff*[f]	Standard error[a] 1	Standard error[a] 2
			A. Nonseparable Model[g]				
1.000	0.9991	1.0800	0.0053	0.0025	−0.0028	0.0048	na
2.000	0.9991	1.0800	0.0053	0.0050	−0.0003	0.0046	na
			B. Abel Model[h]				
1.000	0.9975	3.0025	0.1287	0.0025	−0.1262	0.0182	0.0336
1.500	0.9975	2.9732	0.1259	0.0025	−0.1234	0.0181	0.0374
2.000	0.9976	2.9342	0.1221	0.0025	−0.1196	0.0180	0.0430
2.500	0.9976	2.8807	0.1170	0.0025	−0.1145	0.0179	0.0498
3.000	0.9977	2.8125	0.1105	0.0025	−0.1080	0.0177	0.0572
3.500	0.9977	2.7299	0.1027	0.0025	−0.1002	0.0175	0.0646
6.500	0.9984	1.9280	0.0361	0.0025	−0.0361	0.0125	0.0863
8.500	0.9991	1.0123	0.0053	0.0025	−0.0028	0.0046	0.0061
10.50	0.9999	0.0442	0.0507	0.0025	−0.0482	0.0147	0.1682
13.50	1.0016	−1.9781	0.2827	0.0025	−0.2802	0.0244	0.3827
17.50	1.0046	−5.4836	0.9789	0.0025	−0.9764	0.0301	1.0604
20.00	1.0068	−8.1416	2.6637	0.0025	−2.6612	0.0305	6.7761

[a] All standard errors are constructed using the Newey–West procedure with two lags: (1) standard error on *Diff* ignoring the sampling error in estimating the mean of M, (2) standard error that incorporates the sampling error in estimating the mean of M.
[b] Values higher than 20 imply that the price of the real riskless payoff is outside the arbitrage bound stated in Table 2.
[c] Unconditional mean of M.
[d] $r = (1/E(M) - 1) * 1200$ is the reciprocal of the price of the real riskless bond.
[e] Estimated mean of the growth-optimal return $\tilde{R}_g$.
[f] *Diff* is equal to $E[\ln(1/M)] - E[\ln(\tilde{R}_g)]$.
[g] Time-nonseparable consumption-based model as in Campbell and Cochrane (1995). The mean of M for this model is *assumed* to be consistent with the average of the ex-post real interest rate (see Table 1).
[h] Specification, $M_{t+1} = \lambda_{t+1}^{-\alpha}/\lambda_t^{1-\alpha}$.

product of α and the average growth of consumption, 0.0025 (3% per annum), which must exceed the average of $\ln(\tilde{R}_{gt})$. Table 4 shows that the smallest value of $E[\ln(\tilde{R}_{gt})]$ is 0.0053, which is attained at an average real risk-free bond price of 0.999, a price that corresponds to the sample average of the real T-bill return. Average consumption growth is roughly half the size of the smallest value of the growth-optimal bound and, hence, this specification attains the point estimate of the bound at a risk aversion coefficient of about 2 because $\alpha = E[\ln(\tilde{R}_g)]/E[\ln(\lambda)] = 0.0053/0.0025 = 2.12$.[14] This value is well inside the range that Mehra and Prescott (1985) consider reasonable, and so, one can set α to 2 and search over the free parameters that determine the habit to simultaneously match the real risk-free rate of 1.2%.[15] Moreover, the asymptotic standard errors are so large that the log utility model (i.e., $\alpha = 1$) is not rejected at conventional levels in this sample either. Hence, the growth-optimal bound in this sample has more to say about risk aversion coefficients in these models than it does about the plausibility of the models themselves.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

This model can attain the growth-optimal bound at reasonable values of α whereas the time-separable model cannot because the habit model can match the average risk-free bond price. This extra degree of freedom is obtained by incorporating a mean zero habit variable $\ln(X_{t+1}/X_t)$, which does not affect the mean of $-\ln(M_t)$ but does affect its higher moments. This feature allows these models to provide an explanation of the risk-free rate puzzle and the equity premium puzzle in this two-asset setting with no trading strategies.

In what sense is the growth-optimal bound different from the Hansen–Jagannathan bound for this model? First, note that, unlike the growth-optimal bound (5), the variance of both the postulated proxy and the minimum-variance bound depend on all parameters of the model. Second, as documented by Cochrane and Hansen (1992), the calculation of the variance of the proxy that incorporates these non-separabilities requires an additional assumption regarding the conditional mean of consumption growth but no such assumption is required for the raw-return version of the growth-optimal bound. Finally, the simplicity with which α can be pinned down by the growth-optimal bound is not available in the minimum-variance bound. Of course, we view both of these nonparametric bounds as complementary tools for evaluating the plausibility of asset pricing models.

The Abel (1990) model provides another avenue through which habit formation models can attain the growth-optimal bound. The growth-optimal bound for this model is identical to that for the additively separable log utility model save for the different average riskless bond prices implied by alternative coefficients of relative risk aversion. Perusing Panel B of Table 4 reveals that the asymptotic t-ratios exceed 2.3 in absolute value—that is, a cutoff just below the critical value for the 5% level in a one-tailed test—for values of α between 1 and 2.5. For $\alpha \geq 3$, the large-sample t-statistics are below 1.9 in absolute value and the growth-optimal bound in this model is almost attained numerically for relative-risk-aversion coefficients between 6.5 and 8.5. Hence, the growth-optimal bound in this sample is more informative about plausible risk aversion coefficients for this model than it is about the validity of the model itself.

Although this evidence is favorable to habit formation models, they need not survive an expansion of the asset menu beyond the two-asset case considered here or the incorporation of dynamic trading strategies in computing the growth-optimal bound at comparable risk aversion levels. As noted earlier, altering the analysis in this fashion both increases the growth-optimal bound and raises concerns regarding its small-sample properties. Accordingly, we are currently exploring the behavior of the growth-optimal bound when the asset menu is augmented both by adding additional securities and dynamic trading strategies.

3.5. Alternative Models and Frictions: Some Extensions

Two aspects of the analysis make us think that the growth-optimal bound will provide interesting insights as we seek to apply it to more complicated models and economic settings. First, the clean separability that arises in the specific

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

time-nonseparable models (such as habit formation models) considered here applies in more general models. Second, certain types of market frictions such as borrowing and solvency constraints do not affect the growth-optimal bound, permitting us to apply the existing analysis to representative-agent models with such frictions. We briefly discuss some of these extensions in turn.

The way in which the growth-optimal bound for the internal habit model collapsed to that for the external habit model is symptomatic of a more general probabilistic structure that arises in sufficiently weakly separable models with intertemporal dependencies, including more general habit formation and nonexpected utility models. That is, the marginal utility functional is typically the product of two terms: the direct marginal utility of current consumption (or some service flow perhaps relative to some habit level) and terms involving the ratios of the marginal effect of current consumption on future marginal utility to the direct marginal utility of current consumption. So long as these additional terms are stationary, the expected log of the intertemporal marginal rate of substitution in such a model will be identical to one without such nonseparabilities but with the same-period marginal utilities as in the specific habit formation models discussed earlier. This intuition will fail in models in which the nonseparabilities generate nonstationarities, which we would expect to be the exception rather than the rule under conventional assumptions regarding the stochastic properties of the forcing variables.

Finally, growth-optimal bounds may be useful diagnostic tools to evaluate the effects of market frictions. He and Modest (1995), Luttmer (1996), and Cochrane and Hansen (1992) show that frictions such as solvency and borrowing constraints change the Euler equation equalities to inequalities. Typically, solvency constraints impose the constraint that future wealth of the agent must not fall below a lower bound (which may be state contingent) and borrowing constraints constrain the current consumption choice of the agent to be no greater than his current marketed wealth. However, the Euler inequalities that follow from borrowing and solvency constraints do not change the implications of the growth-optimal bound that $-E[\ln(\tilde{q})] \geq E[\ln(\tilde{R}_g)]$ [see equation (5)]. Frictions, such as short-sale constraints on individual securities also imply Euler inequalities and, thus, leave the growth-optimal bound unaffected save for the requirements that the growth-optimal portfolio weights must respect these constraints as well. Hence, it seems likely that growth-optimal bounds may provide new insights regarding the plausibility of various asset pricing models in the presence of such market frictions.

4. CONCLUSION

Given the premise that frictionless asset markets permit no arbitrage opportunities, we provide a nonparametric necessary condition that any asset pricing model must satisfy in a given asset menu. We show that the absence of arbitrage opportunities implies that the average continuous return on the growth-optimal portfolio provides a lower bound on the average continuous return on the state price density

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

(stochastic discount factor or intertemporal marginal rate of substitution in equilibrium models) implicit in any economic model. Our nonparametric growth-optimal bound is different from the minimum-variance bounds of Hansen and Jagannathan (1991) and provides additional insights regarding aspects of asset market data that are puzzling from the perspective of consumption-based asset pricing models.

The growth-optimal bound captures the economic intuition that empirically reliable estimates of high average compound returns in asset markets should imply correspondingly high average growth rates of the inverse of direct period marginal utility, even in models with time-nonseparabilities. Using the growth-optimal bound, we evaluate the plausibility of time-separable and time-nonseparable consumption-based models. In the context of many popular consumption-based models the growth-optimal bounds translate into bounds for the growth rate of consumption for a given level of risk aversion.

Using a simple asset menu, we document that the time-separable consumption-based model does not satisfy the growth-optimal bounds for plausible risk aversion coefficients. An interesting feature of the growth-optimal bound is that the implications of many time-nonseparable consumption-based models for this bound are similar to those of the time-separable model, although the implications for the average price of the real riskless bond are different across these models. Empirical results suggest that models that incorporate time-nonseparabilities and modest levels of risk aversion can statistically attain the growth-optimal bound and, hence, are not disconfirmed in this sample. However, we suspect that this implication is quite sensitive to the choice of asset payoffs under consideration.

Using alternative asset menus it will be valuable to explore the implications of growth-optimal bounds for various economic models, particularly those that incorporate asset market frictions. To the extent that the implications of the growth-optimal bound remain unaffected either by more general intertemporal nonseparabilities in preferences or by plausible market frictions, the economic insights obtained from this nonparametric tool will both differ from and complement those provided by the Hansen–Jagannathan volatility bounds.

NOTES

1. The growth-optimal portfolio has been studied by Kelley (1956), Latané (1959), Markowitz (1959), Breiman (1960), Hakansson (1971), Roll (1973), Fama and MacBeth (1974), Kraus and Litzenberger (1975), Rubinstein (1976a,b), Grauer (1978), Cox and Huang (1989, 1991), and Long (1990).

2. These realizations are defined implicitly by $\psi(s) = q(s)Pr[s \mid \mathcal{F}]$, where $\psi(s)$ is the state price for state s, $Pr[s \mid \mathcal{F}] = E[I(s) \mid \mathcal{F}]$, and $I(\bullet)$ is the usual indicator function. Note that these probabilities play a perfectly mechanical role; no investor need possess these probability beliefs. That is, the possible values of the state price per unit probability $q(s)$ comprise the events on the relevant sample space and the probabilities $Pr[s \mid \mathcal{F}]$ are the conditional probabilities of these events.

3. Ingersoll (1987) provides a general discussion of all three relations whereas Ross (1978), Harrison and Kreps (1979), and Hansen and Richard (1987) develop the linear pricing rule based on the strong-no-arbitrage condition.

4. When $\mathcal{X}$ does not include a unit payoff, a common occurrence when payoffs are in real terms, Hansen and Jagannathan (1991) also compute versions of p^* and p^+ for different hypothetical values

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

of $E[\tilde{q}]$ which correspond to possible values of the mean of the real risk-free bond price by augmenting $\mathcal{X}$ to include the unit payoff. The assumed values of $E[\tilde{q}]$ and the variance of the minimum variance discount factor characterize the minimum variance frontier for discount factors.

5. Using a quadratic approximation, Cochrane (1991) bounds the variance of pricedividend ratios and provides relations among the variance of price/dividend ratios and forecasts of future dividend growth and mean discount rates.

6. The constrained optimization is given by (4). For any return $\tilde{R}$ included in $\tilde{\boldsymbol{R}}$, the necessary first-order condition is $E[\tilde{R}/\tilde{R}_g] = \delta$, where δ is the multiplier on the constraint $w_g'\iota = 1$. In particular, this pricing restriction also holds for the unit cost return $\tilde{R}_g$, which implies that $E[\tilde{R}_g/\tilde{R}_g] = \delta = 1$. That is, there are only $n - 1$ unknown portfolio weights with n returns.

7. That is, the growth-optimal discount factor satisfies *strong no-arbitrage* by construction. Note also that the growth-optimal return $\tilde{R}_g$ is unique for a given collection of payoffs. The existence of an alternative growth-optimal return $\tilde{G} > 0 \in \mathcal{X}$ would imply both that $E[\tilde{R}_g/\tilde{G}] = 1$ and $E[\tilde{G}/\tilde{R}_g] = 1$. Because the inverse function is strictly convex, both equalities can obtain if and only if $\tilde{R}_g = \tilde{G}$. This uniqueness property is the analog of the uniqueness properly of p^* documented by Hansen and Richard (1987) and Hansen and Jagannathan (1991).

8. A referee pointed out that the Snow's (1991) bounds on higher moments could be substituted into this expression to bound this linear combination. Of course, the log bound provides the tightest lower bound by construction.

9. We have informally examined growth-optimal portfolios using additional assets—a portfolio of government bonds, a small stock portfolio, and a corporate bond portfolio—and two instruments to create managed payoffs—the difference between the 12-month and 3-month T-bill yields and that between the BAA and AAA bonds. Predictably, the addition of securities and trading strategies results in growth-optimal portfolios with much larger ex post continuously compounded average returns with which to confront asset pricing relations. Our small-sample concerns led us to leave detailed analysis of the extent to which these means reflect sharper ex-ante growth-optimal bounds for further research. Additionally, longer postwar return samples roughly double ex post average growth-optimal returns but we confine attention to the smaller sample because comparable data for consumption are not available.

10. For example, with $\alpha = 1$, the standard error of the difference between $\lambda_t^{-\alpha}$ and its mean is only 0.0002, and when α is raised to 4 this standard error rises to 0.0008. Consequently, the large increase in the asymptotic standard error of the growth-optimal bound is due primarily to $E(M_t)$ being close to the arbitrage bound of the real riskless bond price, making $[1 - \hat{w}_g'\iota]$ large.

11. Given the lognormal approximation $\ln(1 + R_f) = \alpha E[\ln(\lambda_t)] - \alpha^2 \operatorname{var}[\ln(\lambda_t)]$, the partial derivative of $\ln(1 + R_f)$ with respect to α is $E[\ln(\lambda_t)] - \alpha \operatorname{var}[\ln(\lambda_t)]$, which is positive for values of α that are less than $\{E[\ln(\lambda_t)]\}/\{\operatorname{var}[\ln(\lambda_t)]\}$ and vice-versa. Hence, an initial increase in risk aversion increases the real interest rate because the effect of positive consumption growth dominates the effects of consumption volatility but the opposite holds for extreme values of α.

12. For example, let $J = 1$, then $h_t = [(c_t + \theta c_{t-1})^{-\alpha} + \beta\theta(c_{t+1} + \theta c_t)^{-\alpha}]/c_t^{-\alpha} = [(1 + \theta/\lambda_t)^{-\alpha} + \beta\theta(\lambda_{t+1} + \theta)^{-\alpha}]$.

13. For example, let $\beta = 1$ and let $\ln(\lambda) = \tau + u$, where u is a normally and independently distributed mean zero random variable; τ and σ_u^2 are chosen to match the consumption growth process observed in the data. Also assume that $\ln(S_{t+1}/S_t) = \delta u$ so that $\ln(S_{t+1}/S_t)$ has mean zero. Given these assumptions, $\ln R_f = \alpha\tau - \alpha^2(1 + \delta)^2\sigma_u^2/2$. For a given τ, σ_u, and α, one can choose δ to match the observed R_f.

14. This point estimate is sensitive to the sample period in question. For example, the average continuously compounded growth-optimal portfolio return is 0.0096 when estimated from December 1953 to December 1993 (that is, extending the sample for returns back before the period for which consumption data are available), roughly doubling the implied relative risk aversion coefficient estimate because 0.0096/0.0025 = 3.84.

15. The Campbell and Cochrane (1995) model provides an illustrative example along these lines.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

REFERENCES

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

Bernardo, A. & O. Ledoit (1996) Gain, Loss and Asset Pricing. Working paper, Anderson School of Management, University of California at Los Angeles.

Breiman, L. (1960) Investment policies for expanding businesses optimal in the long run. *Naval Research Logistics Quarterly* 7, 647–651.

Campbell, J.Y. & J.H. Cochrane (1995) By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. NBER working paper 4995.

Chamberlain, G. & M. Rothschild (1983) Arbitrage and mean-variance analysis on large asset markets. *Econometrica* 51, 1281–1304.

Cochrane, J. (1992) Explaining the variance of price-dividend ratios. *Review of Financial Studies* 5, 243–280.

Cochrane, J.H. & L.P. Hansen (1992) Asset pricing explorations for macroeconomics. In O. Blanchard and S. Fisher (eds.), *NBER Macroeconomics Annual 1992*, pp. 115–165. Cambridge, MA: MIT Press.

Constantinides, G.M. (1990) Habit formation: A possible explanation of the equity premium puzzle. *Journal of Political Economy* 98, 519–543.

Cox, J.C. & C.-f. Huang (1989) Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory* 49, 33–83.

Cox, J.C. & C.-f. Huang (1991) A variational problem arising in financial economics with an application to a portfolio turnpike theorem. *Journal of Mathematical Economics* 20, 465–488.

DeTemple, J.B. & F. Zapatero (1991) Asset prices in an exchange economy with habit formation. *Econometrica* 59, 1633–1657.

Dunn, K.B. & K.J. Singleton (1986) Modelling the term structure of interest rates under habit formation and durability. *Journal of Financial Economics* 17, 27–55.

Fama, E.F. & J.D. MacBeth (1974) Long-term growth in a short-term market. *Journal of Finance* 29, 857–885.

Ferson, W. & G. Constantinides (1991) Habit persistence and durability in aggregate consumption: Empirical tests. *Journal of Financial Economics* 29, 199–240.

Gali, J. (1994) Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit, and Banking* 26, 1–8.

Gallant, A.R., L.P. Hansen & G. Tauchen (1990) Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics* 45, 141–179.

Grauer, R.R. (1978) Generalized two parameter asset pricing models: Some empirical tests. *Journal of Financial Economics* 6, 11–32.

Hakansson, N.H. (1971) multiperiod mean-variance analysis: Toward a general theory of portfolio choice. *Journal of Finance* 26, 857–884.

Hansen, L.P. (1982) Large sample properties of generalized method of moments estimators. *Econometrica* 50, 1029–1054.

Hansen, L.P. & R. Jagannathan (1991) Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.

Hansen, L.P. & R. Jagannathan (1994) Assessing Specification Errors in Stochastic Discount Factor Models. Staff report 167, Federal Reserve Bank of Minneapolis.

Hansen, L.P. & S.F. Richard (1987) The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica* 55, 587–613.

Hansen, L.P. & K.J. Singleton (1982) Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269–1286.

Hansen, L.P., J. Heaton & E. Luttmer (1995) Econometric evaluation of asset pricing models. *Review of Financial Studies* 8, 237–274.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press

Harrison, M.J. & D. Kreps (1979) Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20, 381–408.
He, H. & D. Modest (1995) Market frictions and consumption-based asset pricing. *Journal of Political Economy* 103, 94–117.
Heaton, J. (1995) An empirical investigation of asset pricing with temporally dependent preference specifications. *Econometrica* 63, 353–385.
Ingersoll, J.E. (1987) *Theory of Financial Decision Making*. New Jersey: Rowman & Littlefield.
Kelley, J. R., Jr. (1956) A new interpretation of the information rate. *Bell System Technical Journal* 35, 917–926.
Kocherlakota, N. (1990) On the discount factor in growing economies. *Journal of Monetary Economics* 25, 43–48.
Kraus, A. & R.H. Litzenberger (1975) Market equilibrium in a multiperiod state preference model with logarithmic utility. *Journal of Finance* 30, 1213–1227.
Latané, H.A. (1959) Criteria for choice among risky ventures. *Journal of Political Economy* 67, 144–155.
Long, J.B., Jr. (1990) The numeraire portfolio. *Journal of Financial Economics* 26, 29–69.
Luttmer, E. (1996) Asset pricing in economies with frictions. *Econometrica* 64, 1439–1467.
Markowitz, H.M. (1959) *Portfolio Selection*, Cowles Foundation monograph 16. New York: Wiley.
Mehra, R. & E.C. Prescott (1985) The equity premium: A puzzle. *Journal of Monetary Economics* 15, 145–162.
Roll, R.W. (1973) Evidence on the growth-optimum model. *Journal of Finance* 28, 551–566.
Ross, S.A. (1978) A simple approach to the valuation of risky streams. *Journal of Business* 51, 1–40.
Rubinstein, M. (1976a) The strong case for the generalized logarithmic utility model as the premier model of financial markets. *Journal of Finance* 31, 551–571.
Rubinstein, M. (1976b) The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics and Management Science* 7, 407–425.
Ryder, H.E., Jr. & G.M. Heal (1973) Optimal growth with intertemporally dependent preferences. *Review of Economic Studies* 40, 1–31.
Snow, K.N. (1991) Diagnosing asset pricing models using the distribution of asset returns. *Journal of Finance* 46, 955–983.
Stutzer, M. (1995) A Bayesian approach to diagnostics of asset pricing models. *Journal of Econometrics* 68, 367–397.
Sundaresan, S.M. (1989) Intertemporally dependent preferences and the volatility of consumption and wealth. *Review of Financial Studies* 2, 73–89.

https://doi.org/10.1017/S1365100597003039 Published online by Cambridge University Press