THE JOURNAL OF FINANCE • VOL. XLIX, NO. 1 • MARCH 1994

Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns

STEPHEN G. CECCHETTI, POK-SANG LAM, and NELSON C. MARK*

ABSTRACT

The Euler equations derived from intertemporal asset pricing models, together with the unconditional moments of asset returns, imply a lower bound on the volatility of the intertemporal marginal rate of substitution. This paper develops and implements statistical tests of these lower bound restrictions. While the availability of short time series of consumption data often undermines the ability of these tests to discriminate among different utility functions, we find that the restrictions implied by a number of widely studied financial data sets continue to pose quite a challenge to the current generation of intertemporal asset pricing theories.

RECENT EMPIRICAL RESEARCH ON asset pricing has examined restrictions on the volatility of a representative consumer's intertemporal marginal rate of substitution (IMRS) implied by asset returns data. The pioneering work of Hansen and Jagannathan (1991) shows that the Euler equations derived from a broad range of intertemporal asset pricing models, together with the first two unconditional moments of asset returns, imply a lower bound on the volatility of the IMRS. For an IMRS with a given mean, they derive and compute the minimum standard deviation it must possess. The goal of their work is to restrict the parameter space for a given class of preferences that can be used to understand the dynamics of asset pricing and to evaluate the degree of difficulty encountered by intertemporal asset pricing models in explaining asset returns from various data sets. The purpose of this paper is to develop and implement a statistical procedure for judging whether a particular model of preferences meets Hansen and Jagannathan's lower volatility bound.

The computation of the lower volatility bound has recently developed into a widely used diagnostic tool for assessing the usefulness of a number of classes

* Department of Economics, Ohio State University and NBER; Department of Economics, Ohio State University; and Department of Economics, Ohio State University, respectively. We thank James Bodurtha, In Choi, John Campbell, John Cochrane, Benjamin Friedman, Lars Hansen, John Heaton, Edward Kane, Leonard Santow, Robert Stambaugh, Alan Viard, an anonymous referee, the participants at the NBER Asset Pricing Program Meeting and the seminars at Ohio State, Indiana, and Princeton for helpful comments and suggestions. Cecchetti thanks the National Science Foundation for financial support.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

of preference orderings. For example, Burnside (1990), Epstein and Zin (1991), Hansen and Jagannathan (1991), Heaton (1991), and Ferson and Harvey (1992) all compare the lower volatility bound, computed from stock and bond returns data, with estimates of the mean and standard deviation of the IMRS implied by various utility functions, computed using data on aggregate U.S. consumption. In addition, Snow (1991) examines higher order moments, Bekaert and Hodrick (1992) apply the volatility bound analysis to the study of international equity returns data, and Backus, Gregory, and Telmer (1993) use these methods in an attempt to understand foreign currency returns.[1]

Thus far, researchers have primarily applied this analysis by comparing point estimates of the volatility bound with point estimates of the mean and standard deviation of the IMRS implied by a specific utility function. While the comparison of point estimates may be useful for some purposes, there will be occasions when the investigator will want to employ formal tests of the restrictions that are implied. This paper proposes and carries out such a test.[2]

Our approach is to formulate a procedure that accounts for the two sources of uncertainty that arise in the comparison of the mean and standard deviation of the IMRS implied by a particular model of preferences with the bound that is computed from asset returns data. First, because the volatility bound itself is estimated from the data, it is random. Second, the computation of the mean and standard deviation of the IMRS using a specific utility function relies on estimates of the moments of the consumption process, and so it too is random. As a consequence, a formal statistical evaluation of the restrictions imposed by the implied lower volatility bound can be conducted with a test of whether the difference between two random variables is zero.

We apply our test to four extensively studied data sets, and three popular preference specifications. These include both annual and monthly data on consumption, equity returns, and short-term Treasury debt, as well as data that combine monthly U.S. consumption data with monthly Treasury bill term structure data and with monthly U.S. dollar returns on five major foreign currencies. The utility functions we consider are a one-lag model of consumption durability, a one-lag habit persistence model, and the conventional time-separable constant relative risk aversion (CRRA) model.

An important issue that we address concerns the relative size of the two sources of sampling variation that are present in our test statistic. For high degrees of relative risk aversion, we find that the uncertainty induced by random returns in the estimation of the lower volatility bound, for a given mean of the IMRS, is small relative to the uncertainty in the calculation of the mean and standard deviation of the IMRS based on a model of preferences. Put differently, most of the variation in the comparison of the two

[1] The recent paper by Cochrane and Hansen (1992) provides a survey of work on this topic.

[2] Burnside (1991) has independently devised a set of tests that are similar to the ones discussed in Section I. In addition, Hansen, Heaton, and Luttmer (1993) describe an alternative testing methodology.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

random variables in our test is the result of uncertainty induced by estimation of the mean of the IMRS. This, in turn, is the consequence of uncertainty contained in the consumption data, which is relatively high. We find that the availability of relatively short time series of consumption data—less than 100 years of annual data, and approximately 30 years of monthly data—seriously undermines the ability of tests that use the restrictions implied by the volatility bound to discriminate among different utility functions for the two data sets containing bond and equity returns. The restrictions implied by the Treasury bill term structure data and foreign currency returns data, on the other hand, continue to pose quite a challenge to the current generation of intertemporal asset pricing models.

The remainder of this paper is divided into four sections. In Section I we begin with a review of Hansen and Jagannathan's method for computing the volatility bound from data on asset returns. This is followed by a description of the utility functions we examine, along with a discussion of the stochastic model for consumption. We then show how to compute the IMRS implied by the class of preferences we consider, and derive the statistic used to test whether a model meets the restrictions implied by the volatility bound. Section II reports results of the applications we study. Section III discusses how to impose the restriction that the IMRS is nonnegative and reformulates the testing methodology appropriately. Using this modified procedure, we examine the implication for two of our applications—the annual data on consumption, equity returns, and Treasury debt, and the monthly data on consumption and returns on three-, six-, and nine-month Treasury bills. While the nonnegativity restriction raises the volatility bound of the IMRS, the standard errors grow so much that the test provides less evidence against the null hypothesis. Finally, Section IV provides concluding remarks.

I. A Testing Framework

The purpose of this section is to derive a test to evaluate whether a particular model of preferences is consistent with the restrictions implied by Hansen and Jagannathan's volatility bound. We begin in Section I.*A* with a review of the method used to compute the bound from data on asset returns.

In Section I.*B* we describe the specific utility functions that we examine. These include simple forms of preferences that allow for either consumption durability or habit persistence, as well as the conventional CRRA case. In order to derive the mean and standard deviation of the IMRS implied by the durability and habit persistence specifications, we require knowledge of the consumption process. Section I.*C* describes the stochastic model for consumption that we employ. To this end, we assume that the consumption growth rate follows a random walk in annual data, and a first-order autoregression in monthly data. Section I.*D* then presents the derivation of the mean and standard deviation of the IMRS for the examples we consider.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Finally, Section I.E describes the statistical testing procedure we employ to determine the class of preferences that meet the Hansen and Jagannathan restrictions. Since both the bound itself and the implied volatility of the IMRS for a given utility function depend on data, the comparison of the model to the bound is a test of whether the difference between two random variables equals zero. We exploit this idea, together with standard asymptotic distribution theory, in the derivation of the test.

Throughout this section we ignore one important implication of asset pricing theory—that the IMRS must always be nonnegative. Hansen and Jagannathan show that use of this information can substantially change the location of the volatility bound, and so it can further restrict the set of models that meet the restrictions implied by the bound. We defer discussion of this nonnegativity restriction until Section III, where we present a testing framework in which it is incorporated.

A. The Hansen-Jagannathan Volatility Bound

We begin with a brief description of the derivation of the lower volatility bound on the IMRS first suggested by Hansen and Jagannathan (1991).[3] The starting point is the set of Euler equations implied by intertemporal asset pricing problems. We write these as

$$q_{t-1} = E_{t-1}(v_t x_t) \tag{1}$$

where $E_{t-1}(\cdot)$ is the conditional expectation given information at $t-1$, q_{t-1} is an $(n \times 1)$ vector of asset prices at date $t-1$, x_t is the corresponding vector of date t asset payoffs, and v_t is the intertemporal marginal rate of substitution, which is the discounted ratio of marginal utilities at t and $t-1$. In returns form, q_{t-1} may be a vector of known constants and x_t a vector of gross returns. For example, q_{t-1} might be a vector of ones, so that each asset is defined to command a unit price in return for a stochastic "payoff" equal to its gross return.

To continue, take unconditional expectations of both sides of equation (1), and use the law of iterated expectations, to obtain

$$\mu_q = E(v_t x_t), \tag{2}$$

where μ_q is defined as the unconditional expectation of q_{t-1}, $E(q_{t-1})$. Next, define $\mu_v = E(v_t)$, $\sigma_v^2 = E(v_t - \mu_v)^2$, $\mu_x = E(x_t)$ and $\Sigma_x = E(x_t - \mu_x)(x_t - \mu_x)'$, and then project $(v_t - \mu_v)$, the deviation of the IMRS from its mean, onto $(x_t - \mu_x)$, the deviation of the asset payoffs from their means, to obtain a set of coefficients ρ_v, such that

$$(v_t - \mu_v) = (x_t - \mu_x)'\rho_v + u_t, \tag{3}$$

[3] In addition to the exposition in Hansen and Jagannathan, and the one presented below, there are numerous ways to describe the derivation of the volatility bound. See, for example, Cochrane and Hansen (1992) for another alternative.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

where u_t is the projection error. Using the definitions of the unconditional means and variances, we can write

$$\begin{aligned}\rho_v &= \Sigma_x^{-1}[E(x_t - \mu_x)(v_t - \mu_v)] \\ &= \Sigma_x^{-1}[E(x_t v_t) - \mu_x \mu_v] \\ &= \Sigma_x^{-1}[\mu_q - \mu_x \mu_v],\end{aligned} \tag{4}$$

where the last equality in (4) makes use of the Euler equation (2). Now, using (3) and (4), together with the fact that the projection error u_t is orthogonal to x_t, we can derive the variance of the IMRS. We write this as

$$\sigma_v^2 = (\mu_q - \mu_v \mu_x)'\Sigma_x^{-1}(\mu_q - \mu_v \mu_x) + E(u_t^2). \tag{5}$$

Since $E(u_t^2)$ is nonnegative, it follows that

$$\sigma_v \geq \sigma_x \equiv \left[(\mu_q - \mu_v \mu_x)'\Sigma_x^{-1}(\mu_q - \mu_v \mu_x)\right]^{1/2}. \tag{6}$$

The right-hand side of (6) is the lower volatility bound derived by Hansen and Jagannathan (1991) and we label it σ_x. Provided that μ_q is not a zero-valued vector, the bound is a parabola in (μ_v, σ_v)-space. As Hansen and Jagannathan note, the derivation of this volatility bound may be viewed as the dual to the mean–standard-deviation efficient frontier analysis in the theory of finance, except that there is no guarantee that the returns used to generate the IMRS volatility bound are on the efficient frontier.

A common practice in examining the implications embodied in the lower bound is as follows. First, a point estimate of the bound is computed using point estimates of μ_x and Σ_x from data on asset returns. Next, point estimates of μ_v and σ_v are computed using a particular utility function and consumption data. The investigator then asks which values (if any) of the preference parameters for the utility function result in (μ_v, σ_v) pairs that lie inside the parabola.

As we note in the introduction, this procedure may be useful for addressing certain issues. But the question we ask is whether the mean and standard deviation of the IMRS of the model is "close" to the parabola in a statistical sense. The purpose of the remainder of this section is to describe a method for answering this question.

B. *Preferences*

The main use of the volatility bound is to provide a set of restrictions that allow us to restrict the parameter space for a given class of utility functions. We begin by studying the utility functions examined by Hansen and Jagan-

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

nathan (1991). They assume that the period utility function displays constant relative risk aversion defined over consumption services derived at t, S_t. Expected utility is the discounted expected sum of period utilities, and is written as

$$U_t = E_t \sum_{k=0}^{\infty} \rho^k \frac{S_{t+k}^{1-\gamma} - 1}{(1-\gamma)}, \tag{7}$$

where γ is the coefficient of relative risk aversion, ρ is the discount factor, and consumption services are generated by a simple one-lag model in consumption expenditure, C_t, with coefficient δ:

$$S_t = C_t + \delta C_{t-1}. \tag{8}$$

This utility function includes three cases of interest. When $\delta = 0$, (7) is the familiar CRRA formulation. For $\delta > 0$, (7) implies that consumption purchases contain a durable component of the type studied by both Dunn and Singleton (1986) and Eichenbaum, Hansen, and Singleton (1988). Finally, negative values of δ imply the kind of habit persistence Constantinides (1990) has found useful in explaining the equity premium puzzle.[4]

Using (7) and (8), we can write the IMRS between t and $t+1$ as

$$\text{IMRS}_{t,t+1} = \frac{\rho[S_{t+1}^{-\gamma} + \rho\delta E_{t+1}(S_{t+2}^{-\gamma})]}{S_t^{-\gamma} + \rho\delta E_t(S_{t+1}^{-\gamma})}. \tag{9}$$

We will use (9), together with consumption data, to compute the mean and standard deviation of the IMRS, given values of the preference parameters (ρ, γ, δ). For the familiar CRRA case in which $\delta = 0$, $\text{IMRS}_{t,t+1} = \rho (C_{t+1}/C_t)^{-\gamma}$, we estimate the mean and standard deviation, μ_v and σ_v, nonparametrically from consumption data. For nonzero values of δ we must evaluate the conditional expectations in (9) parametrically. This requires that we specify the stochastic process governing consumption growth.

C. A Stochastic Process for Consumption Growth

Our goal is to examine data at both annual and monthly frequencies. As such, we must select a model for both annual and monthly consumption growth. It is useful to begin with a brief description of the consumption data we study. The annual real consumption series we use is for per capita nondurables plus services. From 1889 to 1928, this series is the data used in Grossman and Shiller (1981), and was provided by Robert Shiller. Beginning in 1929, and continuing through 1987, we use the National Income and

[4] We note that Heaton (1991) examines a model that combines aspects of both durability and habit persistence, while Ferson and Constantinides (1991) provide empirical evidence for this specification of the utility function.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Product Account series for real personal consumption expenditure on nondurables and services. Monthly data are the seasonally adjusted series on real per capita consumption of nondurables and services from April 1964 to December 1988 obtained from CITIBASE.

In order to choose a stochastic model, we first estimate a fourth-order autoregression for the annual data, and a twelfth-order autoregression for the monthly data.[5] The top panel of Table I reports ordinary least squares estimates of these simple autoregressions. The results clearly suggest that the annual data are well approximated by a random walk, and so this is the model that we use. For the monthly data, the coefficient on the first lag of consumption growth is -0.2989 with a t-statistic of 6.3, and a Wald test fails to reject that the second through twelfth coefficients are zero simultaneously. The p-value of this joint test is 0.121. We take these results to suggest that monthly consumption growth can be accurately modeled as an AR(1). The final estimates are reported in the bottom panel of Table I.[6]

D. The Mean and Standard Deviation of the IMRS

Using the stochastic model for consumption growth, we can now compute the mean and standard deviation of the IMRS implied by the preferences described in Section I.*B*. We consider both the case in which the sampling interval for the data and the holding period interval over which returns are computed are the same, and the one in which they are not. For the case of monthly data, this means that we are examining monthly data on both one- and three-month holding period returns, using a stochastic model of consumption that is assumed to be monthly.

We begin with the simpler case in which the holding period interval and the sampling interval coincide. First, write the consumption growth process with autoregressive parameter ω as

$$m_t = \mu_c(1 - \omega) + \omega m_{t-1} + \varepsilon_t, \tag{10}$$

where m_t is the consumption growth rate, defined as $\ln(C_t/C_{t-1})$ and ε_t is an i.i.d. normal random variable with mean zero and variance σ_c^2. Using (10),

[5] We find that our main results are robust to changes in the process for consumption. See footnote 15 below.

[6] We realize that the AR(1) for monthly data does not aggregate to a random walk at an annual frequency. The inconsistency could easily be explained by the fact that the data are time averaged. As He and Modest (1991) point out, the most common method for dealing with this is to assume a process for spot consumption, and derive the statistical implications for time-aggregated consumption. Heaton (1993) examines this problem at length, and shows that use of data averaged over long periods of time reduces the impact of time-aggregation bias. For other discussions of the problems induced by time aggregation, see Grossman, Melino, and Shiller (1987) and Breeden, Gibbons, and Litzenberger (1989).

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table I

Estimated Consumption Processes

Estimated coefficients and standard errors for annual and monthly U.S. consumption growth rate autoregression: $m_t = \Sigma_i \, \omega_i m_{t-i} + \varepsilon_t$, where m_t is the consumption growth rate with mean μ_c, ω_i is the coefficient on the ith lag of consumption growth, and ε_t is an i.i.d. normal innovation with variance σ_c^2. The parameter estimates and their standard errors are taken from generalized method of moments estimates of the parameter vector θ and its covariance matrix, Σ_θ, and are robust to conditional heteroskedasticity.

Panel A. Autoregressions

	Annual Data 1890–1987 (Fourth-Order Autoregression)			Monthly Data 1964:4–1988:12 (Twelfth-Order Autoregression)	
	Coefficient	Standard Error		Coefficient	Standard Error
Constant	0.0185	0.0054	Constant	0.0011	0.0004
ω_1	−0.0972	0.1398	ω_1	−0.2989*	0.0470
ω_2	0.1651	0.1343	ω_2	0.0329	0.0521
ω_3	−0.0670	0.0904	ω_3	0.1130	0.0544
ω_4	−0.0567	0.1085	ω_4	−0.0254	0.0596
			ω_5	0.0185	0.0549
			ω_6	0.0117	0.0497
			ω_7	0.0452	0.0517
			ω_8	0.0801	0.0462
			ω_9	0.0304	0.0531
			ω_{10}	0.1134*	0.0543
			ω_{11}	0.1463*	0.0492
			ω_{12}	0.0119	0.0504

Panel B. Final Estimates

	Annual Data 1890–1987 (Random Walk)		Monthly Data 1964:04–1988:12 (First-Order Autoregression)	
Parameter	Estimate	Standard Error	Estimate	Standard Error
μ_c	0.0172	0.0029	0.0016	0.0002
σ_c^2	0.0012	0.0003	1.9e-5	1.5e-6
ω	—	—	−0.2839	0.0511

Asterisks indicate significance at the 5 percent level.

the IMRS, (9), can be rewritten as

$$\mathrm{IMRS}_{t,t+1} = \mathscr{K}(m_t, m_{t+1})$$

$$= \frac{\rho e^{-\gamma m_t}\left[(e^{m_{t+1}} + \delta)^{-\gamma} + \rho\delta e^{-\gamma m_{t+1}} E_{t+1}(e^{m_{t+2}} + \delta)^{-\gamma}\right]}{(e^{m_t} + \delta)^{-\gamma} + \rho\delta e^{-\gamma m_t} E_t(e^{m_{t+1}} + \delta)^{-\gamma}} \tag{11}$$

where the conditional expectation in (11) is calculated as

$$E_t(e^{m_{t+1}} + \delta)^{-\gamma} = \int_{\underline{\varepsilon}(m_t)}^{\infty} [e^{(\mu_c(1-\omega)+\omega m_t + \varepsilon)} + \delta]^{-\gamma} \Phi_c(\varepsilon)\, d\varepsilon,$$

$$\underline{\varepsilon}(m_t) = \begin{cases} -\infty & \text{if } \delta \geq 0 \\ \ln(-\delta) - \mu_c(1-\omega) - \omega m_t & \text{if } \delta < 0 \end{cases}$$

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

and $\Phi_c(\varepsilon)$ is the normal p.d.f. with mean zero and variance (σ_c^2).[7] It follows from (10) that the unconditional distribution of (m_t, m_{t+1}) is bivariate normal

$$\begin{bmatrix} m_t \\ m_{t+1} \end{bmatrix} \sim N\left[\begin{pmatrix} \mu_c \\ \mu_c \end{pmatrix}, \frac{\sigma_c^2}{1-\omega^2}\begin{pmatrix} 1 & \omega \\ \omega & 1 \end{pmatrix}\right].$$

We denote this distribution $\Phi(m_t, m_{t+1})$. Now, using (11), we compute the mean and standard deviation of the IMRS as[8]

$$\mu_v = \int_{\underline{m}}^{\infty}\int_{\underline{m}}^{\infty} \mathscr{K}(m, m')\Phi(m, m')\, dm\, dm' \tag{12}$$

and

$$\sigma_v^2 = \int_{\underline{m}}^{\infty}\int_{\underline{m}}^{\infty} [\mathscr{K}(m, m') - \mu_v]^2 \Phi(m, m')\, dm\, dm', \tag{13}$$

where

$$\underline{m} = \begin{cases} -\infty & \text{if } \delta \geq 0 \\ \ln(-\delta) & \text{if } \delta < 0 \end{cases}.$$

[7] The consumption process in (10) implies that ε can take on large negative values. When δ is negative, and ε is sufficiently negative, then (11) would require that we calculate the value of a negative number raised to the power $-\gamma$. This problem leads us to put a lower bound on the value of ε when we integrate the expression used to compute the IMRS. This lower bound is labeled $\underline{\varepsilon}(m_t)$. The definition of $E_t(e^{m_{t+1}} + \delta)^{-\gamma}$ in (11) is not entirely consistent with the stochastic process for consumption being correct, since (10) permits ε to take on large negative values with nonzero probability. In practice, this is not a serious problem as the region of the sample space for which $\varepsilon < \underline{\varepsilon}(m_t)$ has negligible probability for all of the cases we consider. For example, let $m^* = \mu_c - 2\sigma_c/\sqrt{1-\omega^2}$ be a value two standard deviations below the mean of m_t, and $z \sim N(0, 1)$. Then using the estimates in Table I and $\delta = -0.5$, $P(\varepsilon < \ln(-\delta) - \mu_c(1-\omega) - \omega m^*)$ is equal to $P(z < -20.51)$ for the annual consumption process, and is equal to $P(z < -160)$ for the monthly consumption process. The quadrature rule we use to compute the discrete approximation to the normal density never strays into the portion of the density where $\varepsilon < \underline{\varepsilon}(m_t)$ so we obtain the same results regardless of whether the lower truncation on ε is imposed.

[8] We introduce $\underline{m}$ in equations (12) and (13) for reasons exactly analogous to the ones that required $\underline{\epsilon}(m_t)$ in equation (11). Again, it makes no difference to our results whether the lower truncation point is imposed because the region of the sample space excluded by the lower truncation has negligible probability. To see this, let z and z' be independent standard normal variates, $a = (\underline{m} - \mu_c)\sqrt{(1-\omega^2)}/\sigma_c$ and $b(z) = (\underline{m} - \mu_c)/\sigma_c - \omega z/\sqrt{(1-\omega^2)}$. Then using the orthogonal transformation, $m = \mu_c + \sigma_c z/\sqrt{(1-\omega^2)}$ and $m' = \mu_c + \sigma_c[\omega z + z'\sqrt{(1-\omega^2)}]/\sqrt{(1-\omega^2)}$. Now, the probability of a realization in the region of the sample space rendered inadmissible by the lower truncation point is $P[z < a, z' < b(z)] = \int_a^{\infty}\int_{b(z)}^{\infty}\phi(z)\phi(z')\, dz'\, dz$ where $\phi(\cdot)$ is the standard normal p.d.f. With $\delta = -0.5$ and the estimates in Table I, we have $a = b(z) = -20.51$ for the annual consumption process. For the monthly process, we have $a = -146.5$ and the lower limit in the first integral is $b(-146.5) = -196.21$.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

When the holding period interval is k periods, for a data-sampling interval of one period, then the relevant IMRS is the one from period t to period $t+k$. Using the consumption process (10), we can write this as

$$\mathrm{IMRS}_{t,t+k} = \Lambda\left(m_t, m_{t+k}, \sum_{i=0}^{k-1} m_{t+i}\right)$$

$$= \rho^k\left(e^{-\gamma\sum_{i=0}^{k-1} m_{t+i}}\right)\left\{\frac{\left[(e^{m_{t+k}}+\delta)^{-\gamma}+\rho\delta e^{-\gamma m_{t+k}}E_{t+k}(e^{m_{t+k+1}}+\delta)^{-\gamma}\right]}{\left[(e^{m_t}+\delta)^{-\gamma}+\rho\delta e^{-\gamma m_t}E_t(e^{m_{t+1}}+\delta)^{-\gamma}\right]}\right\}. \quad (14)$$

The consumption process implies that $\{m_t, m_{t+k}, \sum_{i=0}^{k-1} m_{t+i}\}$ is multivariate normal, given by

$$\begin{bmatrix} m_t \\ m_{t+k} \\ \sum_{i=0}^{k-1} m_{t+i} \end{bmatrix} \sim N\left[\begin{pmatrix} \mu_c \\ \mu_c \\ k\mu_c \end{pmatrix}, \frac{\sigma_c^2}{1-\omega^2}\begin{pmatrix} 1 & \omega^k & \sum_{i=0}^{k-1}\omega^i \\ \omega^k & 1 & \sum_{i=1}^{k}\omega^i \\ \sum_{i=0}^{k-1}\omega^i & \sum_{i=1}^{k}\omega^i & k+2\sum_{i=1}^{k-1}(k-i)\omega^i \end{pmatrix}\right],$$

which we label Γ.

The analogs to (12) and (13) follow as

$$\mu_{vk} = \int_{-\infty}^{\infty}\int_{\underline{m}}^{\infty}\int_{\underline{m}}^{\infty}\Gamma(m, m', m'')\Lambda(m, m', m'')\,dm\,dm'\,dm'' \quad (15)$$

and

$$\sigma_{vk}^2 = \int_{-\infty}^{\infty}\int_{\underline{m}}^{\infty}\int_{\underline{m}}^{\infty}[\Gamma(m, m', m'') - \mu_{vk}]^2\Lambda(m, m', m'')\,dm\,dm'\,dm'' \quad (16)$$

with $\underline{m}$ defined as it is in the simple case above.

In the applications, all of these integrals are evaluated using a 13-point Gauss-Hermite quadrature rule.

E. Testing the Restrictions of the Volatility Bound

We now examine whether the model implied value for the standard deviation of the IMRS is consistent with the bound derived from the asset returns data. This involves asking whether the model value estimated using consumption data, $\hat{\sigma}_v$, is near the bound implied by the asset returns data, using equation (6).

In order to conduct such a comparison, begin by defining ψ as the vector of parameters associated with the stochastic process governing consumption growth. For the AR(1) model of Section I.*C*, $\psi = (\mu_c, \sigma_c, \omega)$. Next, define ϕ as the vector of parameters that characterize the utility function, (ρ, γ, δ).

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Finally, recall that μ_q is the mean vector of the asset prices, and μ_x and Σ_x are the mean vector and covariance matrix of asset payoffs.

We now stack all of the parameters that must be estimated from the data into the vector θ, such that

$$\theta = \begin{pmatrix} \mu_q \\ \mu_x \\ \mathrm{vec}(\Sigma_x) \\ \psi \end{pmatrix},$$

where $\mathrm{vec}(\Sigma_x)$ is the vector obtained by stacking all of the unique elements of the symmetric matrix Σ_x.[9] Now let θ_o be the true value of θ, and $\hat{\theta}$ be a consistent estimator of θ_o such that $\sqrt{T}(\hat{\theta} - \theta_o) \xrightarrow{D} N(0, \Sigma_\theta)$. We presume that we have available a consistent estimator of both θ_o and Σ_θ. In the applications, we compute $\hat{\theta}$ and $\hat{\Sigma}_\theta$ by generalized method of moments using the first two moments of asset returns, the first two moments of consumption growth, and the first-order autocovariance of consumption growth.[10] The covariance matrix $\hat{\Sigma}_\theta$ is the Newey and West (1987) covariance matrix estimator with 11 lags.[11]

Using this notation, we can make explicit the fact that the moments of the IMRS and the volatility bound both depend on the sample. The estimated mean and standard deviation of the model values of the IMRS are $\hat{\mu}_v = \mu_v(\phi;\, \hat{\psi})$ and $\hat{\sigma}_v = \sigma_v(\phi;\, \hat{\psi})$, while the estimated volatility bound is

$$\hat{\sigma}_x = \sigma_x(\phi;\, \hat{\theta}) = \left[\left(\hat{\mu}_q - \mu_v(\phi;\, \hat{\psi})\hat{\mu}_x\right)'\hat{\Sigma}_x^{-1}\left(\hat{\mu}_q - \mu_v(\phi;\, \hat{\psi})\hat{\mu}_x\right)\right]^{1/2}. \quad (17)$$

The comparison of the estimated volatility bound, $\sigma_x(\phi;\, \hat{\theta})$, and the estimated model implied standard deviation of the IMRS, $\sigma_v(\phi;\, \hat{\psi})$, can be carried out by examining the difference

$$\Delta(\phi;\, \hat{\theta}) = \sigma_v(\phi;\, \hat{\psi}) - \sigma_x(\phi;\, \hat{\theta}). \quad (18)$$

[9] For most of our applications, μ_q is simply a vector of known constants, and so in practice it can be omitted from the specification of θ.

[10] The moment conditions used in estimation are

$$\begin{aligned} E[x_t - \mu_x] &= 0 \\ E[\mathrm{vec}(x_t x_t') - (\mathrm{vec}(\Sigma_x) + \mathrm{vec}(\mu_x \mu_x'))] &= 0 \\ E[m_t - \mu_c] &= 0 \\ E\left[m_t^2 - \left(\frac{\sigma_c^2}{1-\omega^2} + \mu_c^2\right)\right] &= 0 \\ E\left[m_t m_{t-1} - \left(\frac{\omega}{1-\omega^2}\sigma_c^2 + \mu_c^2\right)\right] &= 0 \end{aligned}$$

[11] The results do not appear to be sensitive to the number of lags used to compute the robust covariance matrix. For example, there is little change when only three lags are used.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

In order to evaluate whether this difference is large, we require an estimate of the variance of $\Delta(\phi; \hat{\theta})$. This is constructed from the distribution of $\hat{\theta}$.

To proceed, take a mean-value expansion of Δ about θ_o. It follows that

$$\sqrt{T}\left(\Delta(\phi; \hat{\theta}) - \Delta(\phi; \theta_o)\right) \xrightarrow{D} N(0, \sigma_\Delta^2),$$

where

$$\sigma_\Delta^2 = \left(\frac{\partial \Delta}{\partial \theta'}\right)\Bigg|_{\theta_o} \Sigma_\theta \left(\frac{\partial \Delta}{\partial \theta}\right)\Bigg|_{\theta_o}.$$

This is consistently estimated by

$$\hat{\sigma}_\Delta^2 = \left(\frac{\partial \Delta}{\partial \theta'}\right)\Bigg|_{\hat{\theta}} \hat{\Sigma}_\theta \left(\frac{\partial \Delta}{\partial \theta}\right)\Bigg|_{\hat{\theta}}. \tag{19}$$

A test of whether a particular model meets the volatility restriction can now be constructed by testing the null hypothesis, H_o: $\Delta(\phi; \theta_o) \leq 0$. In particular, we compute the ratio $\Delta(\phi; \hat{\theta})/\hat{\sigma}_\Delta$ and look for values of ϕ that make it small. Since this ratio is asymptotically normal, we can use these tests as tools for constructing the regions of the preference parameter space that are not rejected by the volatility bound at various levels of statistical significance. Given that H_o is an inequality, these tests are one-sided, and so appropriate critical values are -1.65 for tests at the five percent level, and -2.33 for tests at the one percent level. Implementation of these tests is the task of Section II.[12]

II. Applications

In this section we test the volatility bound restrictions using four well-known data sets on asset returns. We examine both annual and monthly data on equity returns and short-term Treasury debt in the United States, monthly returns on a portfolio constructed from the U.S. Treasury bills term structure, and monthly U.S. dollar returns on five major foreign currencies. For each data set, we study the three forms of preferences described in Section I.*B*: (1) CRRA, (2) one-lag durability, and (3) one-lag habit persistence. In all cases, we report results for representative values of the preference parameters.[13]

A. Annual U.S. Equity and Short-term Bond Returns: 1890 to 1987

The first returns data set we examine is the one used in Cecchetti, Lam, and Mark (1993) to study the equity premium puzzle. The exact sources of

[12] The procedure we advocate is a Wald test. Obviously, there are other possibilities. For example, since we estimate $\hat{\theta}$ by generalized method of moments, we could add (18) to the list of moments in the estimation, and then test the overidentifying restrictions of the model using the implied J-statistic, as originally described by Hansen (1982). This procedure is similar to the one suggested in Hansen, Heaton, and Luttmer (1993).

[13] The results were computed using FORTRAN programs, and checked with programs written in GAUSS. As a further check on the integrity of our results, both numerical and analytical derivatives were used in most of our computations.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table II

Tests Using Annual U.S. Equity and Short-term Bond Returns

Tests of the volatility bound restrictions using annual data for consumption, returns on the Standard & Poor's index, and one-year Treasury bills (or the equivalent), 1890–1987. For chosen values of the subjective discount factor (ρ), curvature parameter (γ), and lagged consumption parameter (δ), we estimate the mean ($\hat{\mu}_v$) and standard deviation ($\hat{\sigma}_v$) of the IMRS, and the lower volatility bound ($\hat{\sigma}_x$). The test statistic (t-ratio) is constructed under the null hypothesis ($\sigma_v = \sigma_x$).

	$\rho = 0.99$				$\rho = 1.02$			
γ	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio
Panel A. Time-separable Utility: $\delta = 0$								
0	0.990	0.000	0.320	−3.633	1.020	0.000	0.711	−3.073
1	0.973	0.033	0.411	−2.201	1.003	0.034	0.437	−2.563
2	0.958	0.066	0.644	−2.442	0.987	0.068	0.314	−2.537
4	0.931	0.131	1.123	−2.615	0.959	0.135	0.620	−1.420
10	0.881	0.339	2.063	−1.881	0.908	0.349	1.557	−1.343
15	0.874	0.556	2.201	−1.103	0.900	0.573	1.697	−0.753
20	0.900	0.850	1.698	−0.375	0.928	0.876	1.185	−0.137
25	0.967	1.269	0.502	0.276	0.996	1.308	0.358	0.803
30	1.084	1.883	1.889	−0.002	1.116	1.940	2.516	−0.139
Panel B. Time-nonseparable Utility: $\delta = 0.5$ (Durability)								
0	0.990	0.000	0.320	−3.633	1.020	0.000	0.711	−3.073
1	0.973	0.027	0.413	−2.260	1.002	0.028	0.435	−2.607
2	0.957	0.052	0.657	−2.595	0.986	0.054	0.315	−2.562
4	0.928	0.102	1.182	−2.951	0.956	0.105	0.674	−1.717
10	0.861	0.238	2.457	−2.684	0.887	0.246	1.957	−2.136
15	0.825	0.351	3.146	−2.264	0.850	0.362	2.662	−1.881
20	0.805	0.474	3.526	−1.805	0.830	0.490	3.050	−1.509
25	0.800	0.620	3.610	−1.335	0.825	0.641	3.132	−1.098
30	0.812	0.808	3.391	−0.876	0.837	0.836	2.901	−0.687
Panel C. Time-nonseparable Utility: $\delta = -0.5$ (Habit Persistence)								
0	0.990	0.000	0.320	−3.633	1.020	0.000	0.711	−3.073
1	0.979	0.115	0.347	−1.440	1.009	0.121	0.534	−2.067
2	0.982	0.231	0.330	−0.514	1.013	0.244	0.587	−1.225
3	0.997	0.357	0.372	−0.068	1.030	0.377	0.894	−1.032
4	1.028	0.504	0.848	−0.472	1.064	0.535	1.524	−1.167
5	1.077	0.692	1.767	−0.864	1.119	0.744	2.575	−1.329
6	1.154	0.983	3.244	−1.095	1.209	1.105	4.295	−1.397

the data are described in the Appendix to that paper. Briefly, the data set combines the annual consumption data described in Section I.*C* with annual returns on two assets: the Standard and Poor's index supplied by Campbell and Shiller (1987) and one-year U.S. Treasury note yields, or their equivalent. Real returns are computed by deflating nominal values using the Consumer Price Index.

The results of the testing procedure are in Table II. For the purposes of these examples, we have set the discount factor equal to 0.99 and 1.02.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Following the theoretical arguments in Kocherlakota (1990a), who shows that a unique solution to the asset pricing problem exists in economies where the discount factor is greater than one, and the empirical evidence in our earlier paper (Cecchetti, Lam, and Mark (1993)) we include values of ρ that exceed one.[14] The top panel of the table reports the results for the case of time-separable utility ($\delta = 0$), the middle panel includes results for the durability model ($\delta = +0.5$), and the bottom panel presents results for the habit persistence model ($\delta = -0.5$). Each row contains estimates for a particular value of the curvature parameter γ. For each discount factor, we report the model implied values of the mean and standard deviation of the IMRS, $\mu_v(\phi;\ \hat{\psi})$ and $\sigma_v(\phi;\ \hat{\psi})$, the volatility bound evaluated at $\mu_v(\phi;\ \hat{\psi})$, $\sigma_x(\phi;\ \hat{\theta})$, and the t-ratio associated with the comparison of model with the bound, $\Delta(\phi;\ \hat{\theta})/\hat{\sigma}_\Delta$.

The main results are as follows.[15] For the case of time-separable utility and $\rho = 0.99$, the difference between the model and volatility bound is less than 1.65 standard deviations when γ is 10 or higher. With a discount factor of 1.02, values of γ below 4 are consistent with the bound. We contrast this with results that are based solely on the point estimates. Using these data, obtaining point estimates of (μ_v, σ_v) that lie within the volatility bound requires values of γ above 20.[16]

The middle panel of Table II displays results for the one-lag durability model in which $\delta = +0.5$. Again, we report calculations based on γ from 0 to 30, which $\rho = (0.99, 1.02)$. As is well known, for given values of ρ and γ, this

[14] There are a number of ways to understand values of ρ that are greater than one. For example, it can be thought of as a simple, but crude, way of approximating habit formation behavior of the type described in Constantinides (1990). In his model, marginal utility is an increasing function of the level of past consumption. This implies behavior similar to that implied by a discount factor greater than one.

[15] The results reported below are all robust to changes in the specification of the consumption process. Using the annual data, we have examined two additional cases: (1) consumption growth follows a simple AR(1) identical to the one used for monthly data, and (2) consumption growth is governed by the Markov switching model described and estimated in Cecchetti, Lam, and Mark (1990). Using an AR(1), we obtain nearly identical results to the ones reported in Table II. For the Markov switching model, there is no material change in the results either. However, we do find that there is a smaller difference between the time-separable and the one-lag durability cases than is suggested by the results obtained using a random walk. For example, if $\rho = 0.99$ and utility is time separable, the lowest γ not rejected by a t-test at the five percent level is 12 for the random walk case, 14 for the case when consumption growth is an AR(1), and 13 for the Markov switching model. If $\delta = 0.5$, and so preferences exhibit durability, the equivalent values of γ are 22, 23, and 17, respectively.

[16] The relevant parameters are estimated nonparametrically using sample moments of the data in all of our applications of the time-separable model. To investigate the impact of the normality assumption for consumption growth, we also constructed the test statistics by evaluating (12) and (13) setting $\delta = 0$. For the annual data studied in this section, the asymptotic t-ratios are generally larger in absolute value under the normality assumption, thus providing more evidence against the null. The difference between the two methods is small for smaller values of γ. For values of $\gamma = 20$ and above, the t-ratios under normality are as much as 50 percent larger, in absolute value, than those computed nonparametrically. For the monthly data that we examine in Section II.*B* below, there is virtually no difference between the t-ratios computed under the nonparametric and parametric log-normal assumptions.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

model of preferences produces uniformly less variability in the IMRS than the time-separable model. But even so, the one-lag durability model cannot be rejected at the five percent level for $\gamma \geq 20$ when $\rho = 0.99$, and for $\gamma \geq 4$ when $\rho = 1.02$.

Results from the third case we consider, one-lag habit persistence with $\delta = -0.5$, are reported in the bottom panel of Table II. Here we allow for γ ranging from 0 to 6. As one would expect, habit persistence yields substantially more variation in the IMRS for any given value of γ. Consequently, no γ from 1 to 6 is rejected using the 1.65 standard error rule when $\rho = 0.99$. In addition, with $\rho = 1.02$, γ's from two to six are not rejected at the five percent level.

Our results are driven by the fact that the point estimate of the mean of the IMRS for a particular model of preferences, $\hat{\mu}_v(\phi;\, \hat{\psi})$, is very imprecise for larger values of γ.[17] To show this, Table III reports a decomposition of the uncertainty in the comparison of σ_v with σ_x for the time-separable case. We think of this uncertainty as arising from three basic sources. Given the expected value of the IMRS, μ_v, there is uncertainty in both the location of the bound, σ_x, and the standard deviation implied by the model, σ_v. In addition, there is uncertainty induced by the fact that the mean IMRS for the model, μ_v, must be estimated.

For each value of γ, Table III reports a decomposition of the uncertainty in the estimate of $\Delta = (\sigma_v - \sigma_x)$ into its components. The standard error of the estimate of the Hansen-Jagannathan bound for fixed μ_v, $\hat{\sigma}_x$, is in column (2).[18] Columns (3) and (4) report standard errors for $\hat{\sigma}_v$ and $\hat{\Delta}$, again for fixed μ_v. The next three columns of the table, labelled (5), (6), and (7), report the uncertainty in $\hat{\sigma}_x$, $\hat{\sigma}_v$, and $\hat{\Delta}$, that arises solely from randomness in $\hat{\mu}_v$. The final column of the table is our estimate of the "total" standard error in $\hat{\Delta}$, $\hat{\sigma}_\Delta$ computed using the technique described in Section I.*E*.[19]

The results in the table show that for larger values of γ, the main source of uncertainty is the fact that μ_v must be estimated. For example, when γ equals 15, then the estimate of σ_Δ including all sources of uncertainty is 1.49, of which approximately two-thirds can be attributed to the uncertainty arising from the estimation of the mean of the IMRS. The source of the uncertainty in the estimate of μ_v can be linked to the consumption data, since imprecision in estimating the moments of consumption growth lead directly to variance in the estimate of the mean IMRS from the model. Consequently, we conclude that the large standard errors associated with the comparison of the volatility bound to the IMRS moments implied by the

[17] We thank Robert Stambaugh for pointing this out.

[18] Hansen and Jagannathan (1988), an earlier version of the 1991 paper, studies this source of uncertainty.

[19] We note that all of these quantities can be computed by setting various elements in equation (19) to zero. The presence of nonzero covariances implies that entries in the table need not add up to other elements in the same row.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table III

Comparison of Sources of Uncertainty

Standard errors (s.e.) for estimated lower volatility bound ($\hat{\sigma}_x$), IMRS standard deviation ($\hat{\sigma}_v$), and the difference ($\hat{\Delta} = \hat{\sigma}_v - \hat{\sigma}_x$), under alternative assumptions regarding the sources of uncertainty. The utility function is time separable; γ is the coefficient of relative risk aversion and $\rho = 0.99$ is the subjective discount factor. Annual data for consumption, returns on the Standard & Poor's index, and one-year Treasury bills (or the equivalent), 1890–1987.

μ_v:	Fixed	Fixed	Fixed	Random	Random	Random	Random
Returns:	Random	—	Random	Fixed	—	Fixed	Random
γ	s.e.($\hat{\sigma}_x$)	s.e.($\hat{\sigma}_v$)	s.e.($\hat{\Delta}$)	s.e.($\hat{\sigma}_x$)	s.e.($\hat{\sigma}_v$)	s.e.($\hat{\Delta}$)	s.e.($\hat{\Delta}$)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
0	0.088	0.000	0.088	0.000	0.000	0.000	0.088
1	0.151	0.089	0.204	0.037	0.087	0.050	0.172
4	0.184	0.101	0.260	0.226	0.087	0.139	0.380
8	0.257	0.129	0.352	0.505	0.089	0.416	0.724
10	0.286	0.147	0.392	0.670	0.091	0.579	0.917
15	0.305	0.209	0.456	1.190	0.098	1.092	1.492
20	0.240	0.303	0.489	1.911	0.107	1.805	2.265
25	0.160	0.444	0.565	2.346	0.119	2.227	2.774
30	0.432	0.663	0.738	4.472	0.135	4.608	4.010

models of preferences can be traced to the relatively high uncertainty contained in the consumption data.

B. Monthly U.S. Equity and Short-term Bond Returns: 1964 to 1988

We now examine a monthly U.S. data set that combines the consumption data described in Section I.*C* with the return to the Center for Research in Security Prices (CRSP) valued-weighted index of New York Stock Exchange stocks and the one-month holding period return to three-month Treasury bills.[20] The two return series are from the "Fama file" available from CRSP. Real returns are constructed using the implicit price deflator for consumption of nondurables and services. The complete data set extends from April 1964 to December 1988.

The volatility restrictions implied by these monthly bond and equity returns are reasonably difficult to reject under the assumption that investors are risk neutral. However, Hansen and Jagannathan find that a greater challenge to asset pricing models can be posed when conditioning information is incorporated by augmenting the set of asset payoffs, since doing so sub-

[20] It is common to use the one-month Treasury bill for this exercise. But, as discussed in Section II.*C* below, we feel that the one-month data have problems that do not arise at longer maturities.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

stantially raises the volatility bound. Accordingly, in order to test the volatility bound, we multiply each of the original asset prices by the lagged gross return of the assets and the lagged consumption growth rate, thereby expanding the number of assets from 2 to 8. Writing the problem in returns form, denoting the gross return on debt as $r_{1,t}$, the gross return on equity as $r_{2,t}$, and g_t as the gross consumption growth rate we have

$$q_t' = \{1, 1, r_{1,t-1}, r_{2,t-1}, r_{1,t-1}, r_{2,t-1}, g_{t-1}, g_{t-1}\}$$

and

$$x_t' = \{r_{1,t}, r_{2,t}, r_{1,t}r_{1,t-1}, r_{1,t}r_{2,t-1}, r_{2,t}r_{1,t-1}, r_{2,t}r_{2,t-1}, r_{1,t}g_{t-1}, r_{2,t}g_{t-1}\}.$$

Any attempt to meet the restrictions of the volatility bound using monthly data is hampered by the relative smoothness of consumption growth at this frequency. As Hansen and Jagannathan (1991) show, this lack of variation in consumption growth prevents the point estimates of the mean and standard deviation of the IMRS implied by the time-separable ($\delta = 0$) and the one-lag durability ($\delta = +0.5$) models from satisfying the bound unless γ is extremely large. But once sampling variability is taken into account, this is no longer the case. In fact, we are able to find instances in which γ is less than thirty that are not rejected by our test.

Table IV presents the results for the three preference specifications using monthly data on stock and bond returns. Again, the top panel reports findings for the $\delta = 0$ case, while the middle and bottom panels do so for $\delta = +0.5$ and $\delta = -0.5$, respectively. Under habit persistence with $\rho = 1.02$ at an annual rate, there are many values of γ between 2 and 9 that are not rejected at the five percent level—i.e., the t-ratio is below 1.65 in absolute value. In addition, for the time-separable case, $\gamma = 2$ is only marginally rejected with a t-ratio of -1.75. For $\rho = 0.99$, the story is slightly different, with all values of γ being rejected in the time-separable and durability cases, but γ between 4 and 10 not being rejected when there is habit persistence.

These results contrast somewhat with those that are obtained when sampling variability is ignored. For example, Hansen and Jagannathan (1991) conclude that γ must exceed 100 in the CRRA case, and suggest that the monthly data provide a more stringent set of restrictions than the annual data.

C. Monthly U.S. Treasury Bills Term Structure: 1964 to 1988

In choosing the next data set for study, we follow Hansen and Jagannathan and examine the term structure of U.S. Treasury bills data. Specifically we consider a portfolio of three-, six-, and nine-month bills, and use monthly observations on three-month holding period returns constructed from the average of the bid and ask prices. Real returns are computed using the

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table IV

Tests Using Monthly U.S. Equity and Short-term Bond Returns

Tests of the volatility bound restrictions using monthly data for consumption, returns on eight asset returns constructed from the value-weighted NYSE index, and three-month Treasury bills, April 1964 to December 1988. For chosen values of the subjective discount factor (ρ), curvature parameter (γ), and lagged consumption parameter (δ), we estimate the mean ($\hat{\mu}_v$) and standard deviation ($\hat{\sigma}_v$) of the IMRS, and the lower volatility bound ($\hat{\sigma}_x$). The test statistic (t-ratio) is constructed under the null hypothesis ($\sigma_v = \sigma_x$).

	$\rho = 0.99$				$\rho = 1.02$			
γ	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio
Panel A. Time-separable Utility: $\delta = 0$								
0	0.999	0.000	0.158	−2.947	1.002	0.000	1.016	−7.216
1	0.998	0.004	0.662	−4.148	1.000	0.004	0.380	−2.985
2	0.996	0.009	1.300	−5.381	0.998	0.009	0.342	−1.748
4	0.993	0.017	2.565	−5.869	0.995	0.017	1.584	−3.995
10	0.984	0.043	6.167	−5.860	0.986	0.043	5.193	−5.136
15	0.977	0.064	8.941	−5.731	0.979	0.064	7.973	−5.247
20	0.970	0.085	11.510	−5.590	0.973	0.085	10.549	−5.222
25	0.964	0.106	13.878	−5.442	0.967	0.106	12.923	−5.144
30	0.959	0.127	16.049	−5.290	0.961	0.127	15.099	−5.038
Panel B. Time-nonseparable Utility: $\delta = 0.5$ (Durability)								
0	0.999	0.000	0.158	−2.900	1.002	0.000	1.016	−7.068
1	0.998	0.003	0.664	−4.167	1.000	0.003	0.378	−2.984
2	0.996	0.006	1.308	−5.419	0.998	0.006	0.349	−1.790
4	0.993	0.012	2.596	−5.937	0.995	0.012	1.615	−4.075
10	0.983	0.030	6.357	−6.017	0.986	0.030	5.383	−5.305
15	0.976	0.045	9.364	−5.965	0.978	0.045	8.397	−5.493
20	0.968	0.060	12.258	−5.904	0.971	0.060	11.298	−5.548
25	0.961	0.074	15.040	−5.840	0.964	0.074	14.088	−5.554
30	0.955	0.088	17.713	−5.775	0.957	0.088	16.767	−5.534
Panel C. Time-nonseparable Utility: $\delta = -0.5$ (Habit Persistence)								
0	0.999	0.000	0.158	−2.900	1.002	0.000	1.016	−7.068
1	0.998	0.021	0.581	−3.556	1.000	0.021	0.457	−3.349
2	0.997	0.042	0.972	−3.931	0.999	0.042	0.160	−2.166
3	0.996	0.062	1.197	−3.432	0.999	0.063	0.251	−0.765
4	0.996	0.083	1.253	−2.684	0.999	0.084	0.295	−0.597
5	0.996	0.104	1.140	−1.871	0.999	0.105	0.205	−0.281
6	0.997	0.125	0.859	−1.073	1.000	0.126	0.223	−0.217
7	0.998	0.146	0.419	−0.343	1.001	0.147	0.641	−0.634
8	1.000	0.168	0.288	−0.145	1.002	0.168	1.269	−1.127
9	1.002	0.189	1.059	−0.733	1.004	0.190	2.080	−1.596
10	1.004	0.211	2.037	−1.285	1.007	0.212	3.073	−2.022
15	1.024	0.329	9.787	−3.119	1.027	0.330	10.902	−3.468
20	1.057	0.470	23.022	−3.890	1.060	0.473	24.271	−4.067
25	1.109	0.680	43.686	−3.829	1.113	0.683	45.187	−3.927

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

deflator of the nondurable plus service component of consumption. Again, the data are from the "Fama file" supplied by CRSP.[21]

Since we are using monthly observations on three-month holding period returns to compute the bounds, we must compute the model implied values of the mean and standard deviation of the IMRS that mimic this timing. In the notation of Section I, this means that we must calculate $\text{IMRS}_{t,t+3}$, which is just equation (14) with $k = 3$. The model implied values of the moments of the IMRS follow immediately from equations (15) and (16) of Section I.*D*, again setting $k = 3$.[22]

Table V reports the results of using monthly consumption data, together with the three-, six-, and nine-month data on Treasury bills, to test the models of interest. Once again, the three panels of the table refer to the time-separable ($\delta = 0$), one-lag durability ($\delta = +0.5$), and one-lag habit persistence ($\delta = -0.5$) cases, with ρ equal to 0.99 and 1.02.[23] Regardless of the discount factor, all of the time-separable and one-lag durability models reported are rejected at the five percent level, since there are no t-ratios with an absolute value less than 1.65. These results are broadly consistent with those implied by Hansen and Jagannathan's (1991) Figure 6—the Treasury bills data present significant obstacles for asset pricing models.

When we allow preferences to exhibit habit persistence, then the results are somewhat different. Specifically, we find that for γ in excess of 20, the model is not rejected. This is true for both $\rho = 0.99$ and $\rho = 1.02$.

[21] In contrast to Hansen and Jagannathan, we choose not to include the twelve-month Treasury bill, since the currently available CRSP data contain a large number of missing observations—i.e., months in which there was no twelve-month bill outstanding. It has been pointed out to us that the CRSP data set used by Hansen and Jagannathan contained many fewer missing data points in the twelve-month file than the current CRSP releases, because of changes in their sampling procedure. Also, following common practice, we exclude the one-month Treasury bill. Recent work by Luttmer (1991), and Cochrane and Hansen (1992) examines data on one-, three-, six-, and nine-month bills, and finds that the addition of the one-month bill substantially raises the bound. But the one-month bill market is extremely thin, with most trading occurring outside the standard dealer-broker system and transactions costs being substantial—see, for example Stigum (1990, pp. 667ff). As a result, quoted bid-ask spreads are very large, and there is the potential for reported prices to be inaccurate. The methods of both Luttmer and He and Modest, which treat market frictions explicitly and so do not impose the law of one price, are less sensitive to the problems posed by the one-month data.

[22] Our decision to follow common practice in using three-month holding period returns is dictated by our belief that the available data on intermediate term Treasury bills—two, four, five, seven, and eight months to maturity—is of poor quality. We note that if we were to use one-month holding period returns computed from the purchase of a three-, six-, and nine-month bill followed by the sale of a two-, five-, and eight-month bill the following month, then our results are dramatically different. The reason for this is that the one-month holding period returns on each of the three bills have virtually the same mean in the data set, but different variances—it is as if the prices are constructed by simple linear interpolation along the term structure. These data give the false impression that agents are nearly risk neutral. We take this as confirmation of the belief that the only reliable data are those on Treasury bills that are heavily traded.

[23] For the time-separable case, we estimate the moments of the IMRS using the sample moments of $\rho^3(C_{t+3}/C_t)^\gamma$.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table V

Tests Using Monthly Treasury Bill Term Structure

Tests of the volatility bound restrictions using monthly data for consumption, three-, six-, and nine-month Treasury bills, April 1964 to December 1988. For chosen values of the subjective discount factor (ρ), curvature parameter (γ), and lagged-consumption parameter (δ), we estimate the mean ($\hat{\mu}_v$) and standard deviation ($\hat{\sigma}_v$) of the IMRS, and the lower volatility bound ($\hat{\sigma}_x$). The test statistic (t-ratio) is constructed under the null hypothesis ($\sigma_v = \sigma_x$).

	$\rho = 0.99$				$\rho = 1.02$			
γ	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio
Panel A. Time-separable Utility: $\delta = 0$								
0	0.997	0.000	0.253	−1.907	1.005	0.000	1.211	−7.510
1	0.993	0.006	0.513	−2.299	1.000	0.006	0.552	−3.280
2	0.988	0.012	1.157	−3.302	0.995	0.012	0.223	−1.066
4	0.978	0.024	2.445	−3.958	0.985	0.025	1.456	−2.711
10	0.951	0.060	6.132	−4.303	0.958	0.060	5.167	−3.831
15	0.930	0.088	8.993	−4.336	0.937	0.088	8.048	−4.019
20	0.910	0.115	11.669	−4.323	0.917	0.116	10.744	−4.081
25	0.892	0.142	14.168	−4.289	0.898	0.143	13.262	−4.092
30	0.874	0.169	16.498	−4.244	0.881	0.170	15.610	−4.075
Panel B. Time-nonseparable Utility: $\delta = 0.5$ (Durability)								
0	0.997	0.000	0.253	−1.907	1.005	0.000	1.211	−7.510
1	0.993	0.006	0.504	−2.177	1.000	0.006	0.562	−3.213
2	0.988	0.013	1.137	−3.130	0.995	0.013	0.211	−1.016
4	0.978	0.025	2.405	−3.752	0.986	0.025	1.415	−2.523
10	0.952	0.060	6.028	−4.064	0.959	0.061	5.062	−3.595
15	0.931	0.089	8.835	−4.083	0.938	0.089	7.890	−3.768
20	0.912	0.116	11.457	−4.058	0.918	0.117	10.531	−3.818
25	0.894	0.143	13.903	−4.016	0.900	0.144	12.995	−3.819
30	0.877	0.168	16.181	−3.963	0.883	0.170	15.290	−3.795
Panel C. Time-nonseparable Utility: $\delta = -0.5$ (Habit Persistence)								
0	0.997	0.000	0.253	−1.907	1.005	0.000	1.211	−7.510
1	0.993	0.018	0.486	−1.951	1.000	0.018	0.581	−3.014
2	0.988	0.036	1.060	−2.719	0.996	0.036	0.178	−0.913
3	0.984	0.054	1.600	−3.010	0.992	0.054	0.615	−1.318
4	0.981	0.071	2.098	−3.119	0.988	0.072	1.107	−1.826
5	0.977	0.089	2.552	−3.148	0.985	0.090	1.561	−2.098
6	0.974	0.106	2.961	−3.133	0.982	0.107	1.971	−2.242
7	0.972	0.124	3.327	−3.091	0.979	0.125	2.338	−2.310
8	0.969	0.141	3.649	−3.031	0.977	0.142	2.660	−2.331
9	0.967	0.158	3.927	−2.956	0.974	0.160	2.939	−2.318
10	0.965	0.176	4.162	−2.871	0.973	0.178	3.174	−2.280
15	0.962	0.266	4.674	−2.308	0.969	0.268	3.678	−1.850
20	0.966	0.364	4.021	−1.541	0.974	0.367	3.004	−1.139
25	0.981	0.479	2.016	−0.546	0.989	0.484	0.963	−0.174

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

D. Monthly Foreign Currency Returns: 1973 to 1988

The final application we examine uses monthly U.S. dollar speculative returns on five major currencies, together with the monthly consumption data described above. We study spot and one-month forward U.S. dollar prices of the Canadian dollar, the deutsche mark, the French franc, the pound, and the yen. These data are the Friday closing quotations reported in the Harris Bank *Foreign Exchange Weekly Review*. The sample is drawn from those Friday quotations that fall nearest to the end of the calendar month. Again, we calculate real magnitudes using the implicit deflator for consumption of nondurables plus services.

In order to construct the asset portfolios, we begin by defining s_{it} to be the dollar spot price of foreign currency i, and f_{it} to be the one-month dollar forward price of a unit of this currency determined at date t. In determining an investment strategy at date t, an investor will want to go long in the forward foreign currency contract if $(f_{it} - s_{it+1})$ is expected to be positive, and short if $(f_{it} - s_{it+1})$ is expected to be negative.

Let I_{it} be an indicator function such that

$$I_{it} = \begin{cases} 1 & \text{if } E_t(f_{it} - s_{it+1}) > 0 \\ -1 & \text{if } E_t(f_{it} - s_{it+1}) < 0 \end{cases}.$$

The Euler equations implied by this investment strategy can be expressed as

$$0 = E_t\left(v_{t+1} I_{it} \frac{f_{it} - s_{it+1}}{s_{it}} \frac{P_t}{P_{t+1}} \right), \qquad (i = 1, 2, \ldots, n), \tag{20}$$

where P_t is the aggregate U.S. price level.

In notation corresponding to that of Section I, $q_t = 0$ and the ith element of the gross return vector x equals

$$x_{it+1} = I_{it} \frac{f_{it} - s_{it+1}}{s_{it}} \frac{P_t}{P_{t+1}}.$$

Since μ_q is a zero-valued vector, the implied lower volatility bound is a ray from the origin given by

$$\sigma_x = \mu_v \left[\mu_x' \Sigma_x^{-1} \mu_x \right]^{1/2}. \tag{21}$$

Backus, Gregory, and Telmer (1993) investigate the lower volatility bound (21) using univariate data on the five currencies we examine. They evaluate the indicator function by projecting the currency speculative return, $((f_{it} - s_{it+1})/s_{it})$, on the forward premium, $((f_{it} - s_{it})/s_{it})$ and using the fitted values as the estimates of the conditional expectation of the return. The indicator I_{it} is then assigned a value of $+1$ when this fitted value is positive, and -1 when the fitted value is negative.[24]

[24] The projection strategy is defensible on the grounds that the forward premium has proved to be a robust predictor of future currency returns during the modern period of floating exchange rates. See Hodrick (1987) for a survey of theoretical developments and empirical evidence concerning the behavior of foreign currency returns.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

When preferences are given by the one-lag habit persistence model, with a γ of 10, Backus, Gregory, and Telmer find that the point estimate of the volatility of the IMRS, $\hat{\sigma}_v$, is less than one-third of the volatility implied by the foreign currency returns data and the Euler equations. From this, they conclude that these preferences are not capable of explaining the dynamics of foreign currency returns.

Using the Backus, Gregory, and Telmer method for evaluating the indicator function, we examine the portfolio constructed from the five currencies. Again we examine all three preference specifications, $\delta = (0, +0.5, -0.5)$, and two values of the discount factor, $\rho = (0.99, 1.02)$. The results are reported in Table VI. For these data, the volatility bound is substantially higher than it is when we use domestic U.S. stock and bond data. At a mean IMRS value of one, $\mu_v = 1$, the implied lower volatility bound has risen from 0.322 for the monthly stock returns data discussed in Section II.*B*, to 0.402.

It is obvious from the top two panels of Table VI that both the model with time-separable utility and the one with one-lag durability imply much too little variation in the IMRS relative to that implied by the data. For all of the parameter values we consider, both of these cases are easily rejected by our testing procedure, as the t-ratio always exceeds 3.8 in absolute value.

The bottom panel of Table VI reports the results for the one-lag habit persistence model using the data on the five foreign currency returns. Here, values of γ of 15 and higher cannot be rejected at the five percent level.

We cautiously conclude that the restrictions imposed by foreign currencies returns pose quite a challenge to asset pricing models. The source of our prudence is that we have no real sense of how important even small frictions are for these results. As both Luttmer (1991) and He and Modest (1991) suggest, transactions costs, short-sale constraints, and restrictions on borrowing against future labor income can have a very important impact on the height and shape of the volatility bound. It is very possible that once the size of the bid-ask spread is taken into account, then the bound will no longer be as significant as it appears from the results reported in Table VI.

III. Nonnegativity of the IMRS

In this section we generalize the framework described in Section I to incorporate an important implication of asset pricing theory. As Hansen and Jagannathan discuss, the IMRS must be nonnegative or the model implies that assets with a zero probability of a negative payoff will command negative prices. Therefore, it is of interest to reformulate the test of Section I taking into account this information. In the remainder of this section, we begin by discussing the methods we use for computing test statistics. Then in Section III.*B* we report results for the annual data studied in Section II.*A*, and Section III.*C* follows with results for the Treasury bills term structure data studied in Section II.*C*.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table VI

Tests Using Monthly Foreign Currency Returns

Tests of the volatility bound restrictions using monthly data for consumption, and speculative returns on forward foreign exchange on the Canadian dollar, deutsche mark, French franc, pound, and yen, March 1973 to December 1988. For chosen values of the subjective discount factor (ρ), curvature parameter (γ), and lagged-consumption parameter (δ), we estimate the mean ($\hat{\mu}_v$) and standard deviation ($\hat{\sigma}_v$) of the IMRS, and the lower volatility bound ($\hat{\sigma}_x$). The test statistic (t-ratio) is constructed under the null hypothesis ($\sigma_v = \sigma_x$).

	$\rho = 0.99$				$\rho = 1.02$			
γ	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio
Panel A. Time-separable Utility: $\delta = 0$								
0	0.999	0.000	0.402	−5.749	1.002	0.000	0.403	−5.749
1	0.998	0.004	0.401	−5.685	1.000	0.004	0.402	−5.685
2	0.996	0.009	0.401	−5.621	0.999	0.009	0.402	−5.621
4	0.994	0.017	0.400	−5.491	0.996	0.017	0.401	−5.492
10	0.986	0.043	0.397	−5.106	0.989	0.043	0.398	−5.105
15	0.981	0.064	0.394	−4.785	0.983	0.064	0.395	−4.785
20	0.976	0.085	0.392	−4.464	0.978	0.085	0.393	−4.464
25	0.971	0.106	0.391	−4.143	0.974	0.106	0.391	−4.143
30	0.967	0.126	0.389	−3.823	0.969	0.127	0.390	−3.823
Panel B. Time-nonseparable Utility: $\delta = 0.5$ (Durability)								
0	0.999	0.000	0.402	−5.749	1.002	0.000	0.403	−5.749
1	0.998	0.003	0.401	−5.705	1.000	0.003	0.402	−5.705
2	0.996	0.006	0.401	−5.661	0.999	0.006	0.402	−5.661
4	0.994	0.012	0.400	−5.572	0.996	0.012	0.401	−5.572
10	0.986	0.030	0.396	−5.306	0.988	0.030	0.397	−5.306
15	0.980	0.044	0.394	−5.084	0.982	0.044	0.395	−5.084
20	0.974	0.059	0.392	−4.863	0.976	0.059	0.393	−4.863
25	0.968	0.073	0.389	−4.642	0.971	0.073	0.390	−4.642
30	0.963	0.087	0.387	−4.421	0.965	0.087	0.388	−4.421
Panel C. Time-nonseparable Utility: $\delta = -0.5$ (Habit Persistence)								
0	0.999	0.000	0.402	−5.749	1.002	0.000	0.403	−5.749
1	0.998	0.021	0.401	−5.430	1.001	0.021	0.402	−5.429
2	0.997	0.042	0.401	−5.109	1.000	0.042	0.402	−5.107
3	0.997	0.063	0.401	−4.787	1.000	0.064	0.402	−4.785
4	0.997	0.084	0.401	−4.465	1.000	0.085	0.402	−4.461
5	0.998	0.106	0.401	−4.142	1.000	0.106	0.402	−4.138
6	0.999	0.127	0.402	−3.819	1.001	0.128	0.403	−3.814
7	1.000	0.148	0.402	−3.496	1.003	0.149	0.403	−3.490
8	1.002	0.170	0.403	−3.173	1.005	0.171	0.404	−3.166
9	1.005	0.192	0.404	−2.850	1.007	0.193	0.405	−2.843
10	1.008	0.215	0.405	−2.528	1.010	0.216	0.406	−2.519
15	1.030	0.336	0.414	−0.934	1.033	0.338	0.415	−0.922
20	1.067	0.485	0.429	0.544	1.070	0.488	0.430	0.558
25	1.125	0.709	0.452	1.548	1.129	0.714	0.454	1.518

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

A. Methodology

We begin by describing Hansen and Jagannathan's (1988) method for computing point estimates of the volatility bound that impose the nonnegativity restriction. Expressing the problem in returns form, which means that asset prices are all one, the bound they derive is

$$\sigma_x = \left[\lambda^{-1} - \mu_v^2\right]^{1/2}, \tag{22}$$

where μ_v is the mean of the IMRS, λ is defined by

$$\lambda = \min_{w,\alpha} E\{\max[0, w + \alpha' x]\}^2 \tag{23}$$
$$\text{s.t. } w\mu_v + \alpha' l = 1,$$

x is an n-dimensional vector of returns, l is an n-dimensional vector of ones, w is a scalar parameter, and α is an n-dimensional vector of parameters. In the absence of the "max" expression, (23) collapses to the simple case of Section I.*A*.

Hansen and Jagannathan propose estimating σ_x using the sample analog of λ,

$$\hat{\lambda} = \min_{w,\alpha} \frac{1}{T} \sum_{t=1}^{T} \max[0, w + \alpha' x_t]^2 \tag{24}$$
$$\text{s.t. } w\mu_v + \alpha' l = 1.$$

Then an estimate of the volatility bound is

$$\hat{\sigma}_x = \left[\hat{\lambda}^{-1} - \mu_v^2\right]^{1/2}. \tag{25}$$

Following the method in Section I.*E*, we can construct the standard error for the difference between the volatility bound and the model implied value of the standard deviation of the IMRS: $\Delta = \sigma_v - \sigma_x$. As is apparent from (24), $\hat{\lambda}$ depends on the sequence $\{x_t\}$, rather than the sample moments of returns, $\hat{\mu}_x$ and vec($\hat{\Sigma}_x$). This means that we cannot express Δ directly as a function of the estimated moments of consumption and returns, as is done in (18). Nevertheless, we can use the same procedure if we redefine the vector of parameters θ as

$$\theta_\lambda = \begin{bmatrix} \lambda \\ \psi \end{bmatrix}, \tag{26}$$

where λ is defined by (23) and ψ is the parameter vector associated with the consumption growth process, $(\mu_c, \sigma_c, \omega)$.

With this redefinition, we can compute the analog of (19),

$$\hat{\sigma}_\Delta^2 = \left(\frac{\partial \Delta}{\partial \theta_\lambda'}\right)\Bigg|_{\hat{\theta}_\lambda} \hat{\Sigma}_{\theta_\lambda} \left(\frac{\partial \Delta}{\partial \theta_\lambda}\right)\Bigg|_{\hat{\theta}_\lambda}. \tag{27}$$

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Evaluation of (27) requires an estimate of Σ_{θ_λ}, which depends on the covariance of $\hat{\lambda}$ and $\hat{\psi}$, rather than the covariance of the samples moments of returns and $\hat{\psi}$. We can compute an estimate, $\hat{\Sigma}_{\theta_\lambda}$, by stacking the equation that defines λ, (24), together with the moment conditions used to estimate ψ and applying the results of Hansen, Heaton, and Luttmer (1993). The details are described in the Appendix.[25]

B. Annual U.S. Equity and Short-term Bond Returns: 1890 to 1987

The results of including the nonnegativity restriction are presented in Table VII. Here we report estimates of the mean and standard deviation of the IMRS, the volatility bound, and the t-ratio for Δ, using the annual data of Section II.A, time-separable utility and a discount factor of 0.99. For comparison, the table reports values of the bound and the t-ratio both with and without the nonnegativity restriction.

As Hansen and Jagannathan observe, the incorporation of the nonnegativity restriction does sharpen the point estimates of the bound. The effect of the restriction is particularly pronounced when μ_v is small. For example, when μ_v is 0.881, σ_x is 2.063 when nonnegativity is ignored. But with nonnegativity, σ_x rises to 4.749.

This sharpening of the bound necessarily increases the estimate of Δ, the distance between the estimates of σ_v and σ_x. But, as the table shows, the estimated standard error of Δ with nonnegativity increases by so much that, for all but the lowest values of γ we study, the t-ratios are now closer to zero than they were without the restriction. Thus, when taking into account sampling variability, we obtain the surprising result that imposing the nonnegativity constraint raises the volatility bound while at the same time yielding less evidence against the null hypothesis.

We conjecture that one reason for this result is that the added uncertainty driving our results comes from the truncation in the computation of λ. This can be seen from the "max" in (23). As μ_v decreases from the point at which σ_x is minimized (the bottom of the parabola) the truncation caused by the "max" function implies that more and more information is thrown out in the calculation of $\hat{\lambda}$. As less data are used to compute the bound, the precision of the bound deteriorates.

C. Monthly U.S. Treasury Bills Term Structure: 1964 to 1988

We have also examined the importance of the nonnegativity constraint using the same Treasury bills data set studied in Section II.C. The results are presented in Panel B of Table VII. Again, we report estimates both with and

[25] In our working paper (Cecchetti, Lam, and Mark (1992)) we present an alternative, and less complex, method for incorporating the nonnegativity restriction. There we assume that returns are jointly normally distributed, and show how to construct a parametric estimate of the test statistic which imposes the nonnegativity of the IMRS. The results from the parametric method are both qualitatively and quantitatively similar to those obtained using the nonparametric method described in the text.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

Table VII

Tests Incorporating the Nonnegativity of the IMRS

Tests of the volatility bound restrictions that incorporate nonnegativity of the IMRS, using annual data for consumption, and either returns on the Standard & Poor's index, and one-year Treasury bills (or the equivalent), 1890 to 1987; or three-, six-, and nine-month Treasury bills, April 1964 to December 1988. For a subjective discount factor (ρ) of 0.99, time-separable utility, and various values of the curvature parameter (γ), we estimate the mean ($\hat{\mu}_v$) and standard deviation ($\hat{\sigma}_v$) of the IMRS, and the lower volatility bound ($\hat{\sigma}_x$). The test statistic (t-ratio) is constructed under the null hypothesis ($\sigma_v = \sigma_x$).

			Without Nonnegativity		With Nonnegativity	
γ	$\hat{\mu}_v$	$\hat{\sigma}_v$	$\hat{\sigma}_x$	t-Ratio	$\hat{\sigma}_x$	t-Ratio
Annual Data on Equity Returns, 1890–1987						
0	0.990	0.000	0.320	−3.633	0.320	−3.633
1	0.973	0.033	0.411	−2.201	0.411	−2.203
2	0.958	0.066	0.644	−2.441	0.675	−1.975
4	0.931	0.131	1.123	−2.615	1.500	−1.753
10	0.881	0.339	2.063	−1.880	4.749	−0.967
15	0.874	0.556	2.201	−1.103	5.613	−0.473
20	0.900	0.850	1.698	−0.375	3.033	−0.279
25	0.967	1.269	0.502	0.277	0.505	0.261
30	1.084	1.883	1.889	−0.002	2.107	−0.041
Treasury Bills Term Structure: Monthly, 1964–1988						
0.0	0.997	0.000	0.513	−5.377	0.557	−4.347
0.5	0.995	0.003	0.607	−4.225	0.633	−3.986
1.0	0.993	0.006	0.833	−4.004	0.870	−3.633
1.5	0.990	0.009	1.120	−4.122	1.208	−3.412
2.0	0.988	0.012	1.424	−4.248	1.637	−3.124
2.5	0.985	0.015	1.736	−4.348	2.180	−2.814
3.0	0.983	0.018	2.052	−4.419	2.859	−2.580
3.5	0.981	0.021	2.369	−4.480	3.725	−2.329
4.0	0.978	0.024	2.686	−4.554	4.971	−1.902
4.5	0.976	0.027	3.001	−4.558	9.596	−0.588

without the restriction using time-separable utility and a discount factor of 0.99. The table presents results for γ ranging from 0 to 4.5 for the following reason: As γ increases, the estimated IMRS decreases, which increases the estimated volatility bound. Using these data, and in cases in which γ exceeds 5, the result is that the estimated volatility bound is infinite.

The incorporation of the nonnegativity restriction into the calculations using the Treasury bills term structure sharpens the volatility bound considerably. For example, when γ is 4.5, the $\hat{\sigma}_x$ is three times higher with nonnegativity than it is without. Even so, we again observe that the inclusion of the restriction increases the standard error of Δ enough so that we actually have less evidence against the null hypothesis. Note that when nonnegativity is considered, $\gamma = 4.5$ cannot be rejected at the five percent level, while without the restriction it could be rejected at the one percent level.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

IV. Conclusion

This paper has developed and implemented a procedure for testing the restrictions implied by Hansen and Jagannathan's (1991) lower volatility bound for the intertemporal marginal rate of substitution. Our approach allows us to evaluate whether the standard deviation of the IMRS implied by a particular model of preferences is consistent with the bound derived from asset returns data. The result is a statistical test that can be used to formally reject some models.

Previous investigators have concluded that the restrictions implied by the bound allow rejection of many commonly used utility functions for reasonable parameter values. But their methods involve comparison of the point estimates of the bound with the IMRS volatility derived using different utility functions. In contrast to these results, we find that by taking explicit account of sampling variability, the restrictions implied by the bound do not allow rejection of a number of models with reasonable parameter values. In particular, using annual data on equity and bond returns in the United States over the last century, we find that the constant relative risk aversion utility is consistent with the bound when the discount factor is 0.99 and the CRRA coefficient is 6 or higher. From this we conclude that the failure of some models is not nearly as extreme as the point estimates would suggest.

We go on to examine three additional data sets, and discover that we can find a set of preferences that are not rejected by the restrictions implied by the volatility bounds. This is true of: (1) monthly data on stock prices and Treasury debt, where we find that time-separable utility with a discount factor of 1.02 and a CRRA coefficient near two is marginally consistent with the data; (2) monthly Treasury bills term structure data, where we find that one-lag habit persistence utility with a discount factor of 1.02 and a CRRA coefficient of twenty is consistent with the data; and (3) data on returns to five foreign currencies, where we find that preferences exhibiting habit persistence with a discount factor of 0.99 and curvature parameter of fifteen is consistent with the data.

We also examine the importance of explicitly considering the fact that the IMRS must be nonnegative. While the incorporation of the information in this restriction does sharpen the volatility bound, as Hansen and Jagannathan originally found, our results suggest that the uncertainty associated with the location of the bound grows so rapidly as to make it less conclusive than tests that ignore the restriction.

Appendix

This appendix uses the results of Hansen, Heaton, and Luttmer (1993) to construct the covariance matrix for θ_λ, which is required to estimate the difference between the standard deviation of the IMRS and the Hansen-Jagannathan volatility bound imposing the nonnegativity restriction, as dis-

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

cussed in Section III.*A*. For simplicity, we discuss the case of time-separable utility.

Assume that consumption growth is stationary and ergodic. It follows from Hansen (1982) that we can consistently estimate the first two moments of the IMRS by their sample analogs

$$\hat{\mu}_1 = \Sigma_T \rho \left(\frac{C_{t+1}}{C_t} \right)^{\gamma} \tag{A1}$$

$$\hat{\mu}_2 = \Sigma_T \rho^2 \left(\frac{C_{t+1}}{C_t} \right)^{2\gamma}, \tag{A2}$$

where, for convenience, we let $\Sigma_T X_t \equiv 1/T \Sigma_{t=1}^{T} X_t$.

Next, consider the definition of λ from the problem in Section III.*A*, equation (23). Let the vector of returns on the n assets be given by x, then substitute the constraint into the minimization problem. Following closely the notation of Hansen, Heaton, and Luttmer, we write the result as

$$\lambda = E\left\{\tilde{\phi}[\alpha^*(\mu_1), \mu_1]\right\}, \tag{A3}$$

where $\tilde{\phi} = \{\max[0,(1 - \alpha^*(\mu_1)'l)/\mu_1 + \alpha^*(\mu_1)'x]\}^2$, l is a vector of ones, and $\alpha^*(\mu_1) = [\alpha_1^*(\mu_1), \alpha_2^*(\mu_1), \ldots, \alpha_n^*(\mu_1)]'$, the vector of minimizers of the problem in equation (23) for a given μ_1. The sample analog for λ is then

$$\hat{\lambda} = \Sigma_T\left\{\tilde{\phi}[\alpha^*(\hat{\mu}_1), \hat{\mu}_1]\right\}, \tag{A4}$$

where $a^*(\hat{\mu}_1)$ is the vector of minimizers of the problem in equation (24), for a given $\hat{\mu}_1$.

Assuming that $\hat{\lambda}$ is continuously differentiable in $\hat{\mu}_1$, we expand $\sqrt{T}(\hat{\lambda} - \lambda)$ into three terms:

$$\begin{aligned} \sqrt{T}(\hat{\lambda} - \lambda) = {} & \sqrt{T}\left(\Sigma_T\left\{\tilde{\phi}[a^*(\mu_1), \mu_1] - \tilde{\phi}[\alpha^*(\mu_1), \mu_1]\right\}\right) \\ & + \sqrt{T}\left(\Sigma_T\left\{\tilde{\phi}[\alpha^*(\mu_1), \mu_1] - E\left(\tilde{\phi}[\alpha^*(\mu_1), \mu_1]\right)\right\}\right) \\ & + \hat{\lambda}'(\tilde{\mu}_1)\sqrt{T}(\hat{\mu}_1 - \mu_1), \end{aligned} \tag{A5}$$

where $\hat{\lambda}'$ is the derivative of $\hat{\lambda}$ with respect to μ_1, evaluated at a value between μ_1 and $\hat{\mu}_1$, $\tilde{\mu}_1$. We note that it is the final term in (A5) that differentiates our problem from that of Hansen, Heaton, and Luttmer, as we consider explicitly the sampling variability associated with the estimation of μ_1.

Two facts established by Hansen, Heaton, and Luttmer are useful for analyzing (A5). First, the probability limit of the first term is zero. Second, $\hat{\lambda}$ is a consistent estimator for λ for any given μ_1. Since both $\hat{\lambda}$ and $\hat{\mu}_1$ are consistent, it follows that $\hat{\lambda}'(\tilde{\mu}_1)$ converges in probability to $\lambda'(\mu_1)$. Using

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

these results, we can then rewrite equation (A5) as

$$\sqrt{T}(\hat{\lambda}-\lambda) \simeq \sqrt{T}\Big(\Sigma_T\Big\{\tilde{\phi}[\alpha^*(\mu_1),\mu_1] - E\Big(\tilde{\phi}[\alpha^*(\mu_1),\mu_1]\Big)\Big\}\Big) + \lambda'(\mu_1)\sqrt{T}(\hat{\mu}_1-\mu_1). \tag{A6}$$

Next, consider the vector of estimates of the objects of interest, $\hat{\theta}_\lambda = [\hat{\lambda}, \hat{\mu}_1, \hat{\mu}_2]'$. It follows from (A1), (A2), and (A6) that

$$\sqrt{T}\left(\hat{\theta}_\lambda - \theta_\lambda\right) \simeq D\left[\sqrt{T}\Sigma_T f\right], \tag{A7}$$

where

$$D = \begin{bmatrix} 1 & \lambda'(\mu_1) & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and

$$f = \begin{bmatrix} \tilde{\phi}[\alpha^*(\mu_1),\mu_1] - \lambda \\ \rho\left(\dfrac{C_{t+1}}{C_t}\right)^{\gamma} - \mu_1 \\ \rho^2\left(\dfrac{C_{t+1}}{C_t}\right)^{2\gamma} - \mu_2 \end{bmatrix}.$$

Assuming that returns and consumption growth are jointly stationary and ergodic, the results in Hansen (1982) imply that $[\sqrt{T}\Sigma_T f]$ converges in distribution to a normal with mean zero and covariance matrix Ω. It then follows immediately that $\sqrt{T}(\hat{\theta}_\lambda - \theta_\lambda)$ converges in distribution to $N(0, D\Omega D')$. Since we wish to conduct inference with respect to $\hat{\theta}_\lambda$, we need to obtain an estimate of $D\Omega D'$. We compute this from $D_T\Omega_T D_T'$, where D_T is an estimator of D which uses $\hat{\lambda}'(\hat{\mu}_1)$ to estimate $\lambda'(\mu_1)$, and Ω_T is the Newey-West ($m = 11$) estimator of the spectral density at frequency zero of f. This estimated covariance matrix corresponds to $\hat{\Sigma}_{\theta_\lambda}$ in equation (27).

REFERENCES

Backus, David K., Allan W. Gregory, and Chris I. Telmer, 1993, Accounting for forward rates in markets for foreign currency, *Journal of Finance* 48, 1887–1908.

Bekaert, Geert, and Robert J. Hodrick, 1992, Characterizing predictable components in excess returns on equity and in foreign exchange markets, *Journal of Finance* 47, 467–510.

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

Burnside, A. Craig, 1990, The comovement of asset returns with time non-separable preferences, Mimeo, Department of Economics, Queens University.

———, 1991, Hansen-Jagannathan bounds as classical tests of asset pricing models, Mimeo, Department of Economics, Queens University.

Campbell, John Y., and Robert J. Shiller, 1987, Cointegration and tests of the present value model, *Journal of Political Economy* 95, 1062–1088.

Cecchetti, Stephen G., Pok-sang Lam, and Nelson C. Mark, 1990, Mean reversion in equilibrium asset prices, *American Economic Review* 80, 398–418.

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms

———, 1992, Testing volatility restrictions on intertemporal marginal rates of substitution implied by Euler equations and asset returns, N.B.E.R. Technical Working Paper No. 124.

———, 1993, The equity premium and the risk-free rate: Matching the moments, *Journal of Monetary Economics* 31, 21–46.

Cochrane, John H., and Lars Peter Hansen, 1992, Asset pricing lessons for macroeconomics, in O. J. Blanchard and S. Fischer, eds.: *NBER Macroeconomics Annual* vol. 7, (M.I.T. Press, Cambridge, Mass.), pp. 115–164.

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

Dunn, Kenneth, and Kenneth J. Singleton, 1986, Modeling the term structure of interest rates under nonseparable utility and durability of goods, *Journal of Financial Economics* 41, 333–355.

Eichenbaum, Martin S., Lars Peter Hansen, and Kenneth J. Singleton, 1988, A time series analysis of representative agent models of consumption and leisure choice under uncertainty, *Quarterly Journal of Economics* 103, 51–78.

Epstein, Larry G., and Stanley E. Zin, 1991, The independence axiom and asset returns, Mimeo, G.S.I.A., Carnegie-Mellon University.

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

Grossman, Sanford J., and Robert J. Shiller, 1981, The determinants of the variability of stock market prices, *American Economic Review* 71, 222–227.

Grossman, Sanford J., Angelo Melino, and Robert J. Shiller, 1987, Estimating the continuous-time consumption-based asset pricing model, *Journal of Business and Economic Statistics* 5, 315–328.

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

———, John Heaton, and Erzo G. J. Luttmer, 1993, Econometric evaluation of intertemporal asset pricing models using volatility bounds, Mimeo, Department of Economics, University of Chicago.

Hansen, Lars Peter, and Ravi Jagannathan, 1988, Restrictions on intertemporal marginal rates of substitution implied by asset returns, Mimeo, Department of Economics, University of Chicago.

———, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 91, 225–262.

He, Hua, and David M. Modest, 1991, Market frictions and consumption-based asset pricing, Mimeo, Haas School of Business, University of California at Berkeley.

Heaton, John, 1991, An empirical investigation of asset pricing with temporally dependent preference specifications, Mimeo, Department of Economics, M.I.T.

———, 1993, The interaction between time-nonseparable preferences and time aggregation, *Econometrica* 61, 353–386.

Hodrick, Robert J., 1987, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets* (Harwood Academic Publishers, New York).

Kocherlakota, Narayana R., 1990a, A note on the "discount" factor in growth economies, *Journal of Monetary Economics* 25, 43–48.

———, 1990b, On tests of representative consumer asset pricing models, *Journal of Monetary Economics* 26, 285–304.

Luttmer, Erzo G. J., 1991, Asset pricing in economies with frictions, Mimeo, Department of Economics, University of Chicago.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Snow, Karl N., 1991, Diagnosing asset pricing models using the distribution of asset returns, *Journal of Finance* 46, 955–984.

Stigum, Marcia L., 1990, *The Money Market*, 3rd ed. (Richard Irwin Inc., New York).

This content downloaded from 129.74.250.206 on Tue, 16 Jan 2018 19:29:26 UTC
All use subject to http://about.jstor.org/terms