

NBER WORKING PAPER SERIES

STOCHASTIC DISCOUNT FACTOR BOUNDS WITH CONDITIONING INFORMATION

Wayne E. Ferson
Andrew F. Siegel

Working Paper 8789
<http://www.nber.org/papers/w8789>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2002

Wayne E. Ferson is the Pigott-PACCAR Professor of Finance, University of Washington and a Research Associate, National Bureau of Economic Research [phone: (206)543-1843, fax: 685-9392, <http://faculty.washington.edu/apmodels>]. Andrew F. Siegel is the Grant I. Butterbaugh Professor of Finance and Management Science and Adjunct Professor of Statistics, University of Washington [phone: (206)543-4476, fax: 685-9392, e-mail asiegel@u.washington.edu]. We are grateful to Geert Bekaert, Ravi Jagannathan, Jay Shanken, Guofu Zhou, and an anonymous referee for helpful comments. This paper has also benefited from seminars at the 1998 European Finance Association Meetings, the July, 1998 National Bureau of Economic Research Asset Pricing Program, the January, 1999 American Finance Association, and at the following Universities: Arizona State, London Business School, Louisiana State, Miami, Notre Dame, South Carolina, Washington University, and the University of Washington. The paper was previously titled *Optimal Moment Restrictions on Stochastic Discount Factors*. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by Wayne E. Ferson and Andrew F. Siegel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Stochastic Discount Factor Bounds with Conditioning Information
Wayne E. Ferson and Andrew F. Siegel
NBER Working Paper No. 8789
February 2002
JEL No. G12, C31, D51

ABSTRACT

Hansen and Jagannathan (HJ, 1991) describe restrictions on the volatility of stochastic discount factors (SDFs) that price a given set of asset returns. This paper compares the sampling properties of different versions of HJ bounds that use conditioning information in the form of a given set of lagged instruments. HJ describe one way to use conditioning information. Their approach is to multiply the original returns by the lagged variables, and much of the asset pricing literature to date has followed this multiplicative approach. We also study two versions of optimized HJ bounds with conditioning information. One is from Gallant, Hansen and Tauchen (1990) and the second is based on the unconditionally-efficient portfolios derived in Ferson and Siegel (2000). We document finite-sample biases in the HJ bounds, where the biased bounds reject asset-pricing models too often. We provide useful correction factors for the bias. We also evaluate the asymptotic standard errors for the HJ bounds, from Hansen, Heaton and Luttmer (1995).

Wayne E. Ferson
Department of Finance
Box 353200
University of Washington
Seattle, WA 98195
and NBER

Andrew F. Siegel
Department of Finance
Box 353200
University of Washington
Seattle, WA 98195
asiegel@u.washington.edu

1. Introduction

Most asset pricing models can be represented in the form of a fundamental valuation equation:

$$E(m_t R_t | Z_{t-1}) = e \quad (1)$$

where R is a vector of gross (unity plus rate of) returns on traded assets, Z_{t-1} is a vector of instruments in the public information set at time $t-1$ and e is a vector of ones. The standard asset pricing models in finance specify the form of the random variable, m_t , the *stochastic discount factor* (see the review by Ferson, 1995). The elements of the vector $m_t R_t$ may be viewed as “risk adjusted” gross returns. The returns are risk adjusted by “discounting” them, or multiplying by m_t , to arrive at the “present value” per dollar invested, equal to one dollar. A stochastic discount factor is said to “price” the assets in R if Equation (1) is satisfied.

Hansen and Jagannathan (HJ, 1991) derive lower bounds for the variance of any stochastic discount factor which satisfies the fundamental valuation Equation (1); such bounds may be used as a prior diagnostic. If a candidate for m_t , corresponding to a particular theory, fails to satisfy the HJ bounds, then it can not satisfy the Equation (1).

Burnside (1994) describes classical hypothesis tests based on the distance between a stochastic discount factor (SDF) and the HJ bounds. He evaluates the sampling properties of such tests with a Monte Carlo simulation of the consumption-based model from Lucas (1978). Tierens (1993) extends the simulation evidence to the Epstein-Zin (1991) model. Both studies find that the sample SDF plots outside the sample HJ bounds too often when the model is true. However,

both studies limit their attention to cases where there is no conditioning information, so the lagged instrument, Z_{t-1} , is a constant.

This paper focuses on the use of conditioning information in the HJ bounds. For a given choice of lagged variables, Z_{t-1} , there are several ways to implement the bounds, but no previous study compares the properties of the various approaches. Given Burnside's (1994) and Tieren's (1993) evidence of biases in models with no conditioning information, an analysis of the sampling properties of bounds with conditioning information is important.

We evaluate the finite-sample properties of HJ bounds with three approaches to conditioning information: (1) the multiplicative approach suggested by HJ; (2) the optimal bounds of Gallant, Hansen and Tauchen (1990); and (3) the efficient portfolio bounds, based on the unconditionally efficient portfolios derived by Ferson and Siegel (2000). We also evaluate asymptotic standard errors for the HJ bounds, derived by Hansen, Heaton and Luttmer (HHL, 1995). Our results show that the use of conditioning information in the bounds is important, and the way in which information is used is important. When sampling error is accounted for, bounds that use no conditioning information have little economic content, while bounds that use the conditioning information efficiently do have economic content. The optimized bounds have significantly more economic content than the multiplicative bounds, once we adjust for sampling errors.

We document finite-sample biases in HJ bounds with conditioning information. The biases imply that the bounds reject asset-pricing models too often. We argue that the magnitudes of the biases are economically significant. The bias is the largest in the multiplicative approach and the smallest for the efficient portfolio bounds.

We provide corrections for the finite-sample bias for each version of the HJ bounds. In addition to correcting for bias in the location of the bounds, our adjustment produces bounds with smaller standard errors. Our simulations show that the adjustments approximately remove the bias in multiplicative and efficient portfolio bounds, but a bias remains in the Gallant, Hansen and Tauchen (1990) bounds.

The paper is organized as follows. In section 2 we review the original HJ bounds. Bounds with conditioning information are discussed in Section 3. We present the adjustments for finite sample bias in Section 4. Section 5 describes the data that we use in our empirical examples and Section 6 provides the examples. Section 7 describes a simulation study into the properties of the various methods for computing the HJ bounds. Section 8 considers the effects of nonnormality and heteroskedasticity in the data. Section 9 evaluates the asymptotic standard errors of HHL. Section 10 offers a summary and concluding remarks.

2. The Hansen-Jagannathan Bounds

We first consider the special case where the conditioning information is a constant, so the expectations in (1) are unconditional.

Assume that the random column n -vector R of the assets' gross returns has mean $E(R)=\mu$ and covariance matrix Σ . When there is no conditioning information a stochastic discount factor is defined as any random variable m such that $E(mR)=e$.

Proposition 1 (Hansen and Jagannathan, 1991). The stochastic discount factor m with minimum variance for its expectation $E(m)$ is given by

$$m = E(m) + [e - E(m)\mu]' \Sigma^{-1} (R - \mu). \quad (2)$$

And the variance of m is

$$\sigma_m^2 = [e - E(m)\mu]' \Sigma^{-1} [e - E(m)\mu]. \quad (3)$$

The proof is provided in Hansen and Jagannathan (1991).

Hansen and Jagannathan (1991) show that their bound is related to the maximum Sharpe ratio that can be obtained by a portfolio of the assets under consideration. The Sharpe ratio is defined as the ratio of the expected excess return to the standard deviation of the portfolio return. If the vector of assets' expected excess returns is $\mu - E(m)^{-1}e$ and Σ is the covariance matrix, the square of the maximum Sharpe ratio is $[\mu - E(m)^{-1}e]' \Sigma^{-1} [\mu - E(m)^{-1}e]$. Thus, from Equation (3) the lower bound on the variance of stochastic discount factors is the maximum squared Sharpe ratio multiplied by $[E(m)]^2$.

3. Bounds with Conditioning Information

Recent papers refine and extend the HJ bounds in several directions.² This paper focuses on the use of given lagged variables, Z_t , to refine the bounds. To understand how such conditioning

² HJ (1991) show how restricting $m > 0$ can refine the bounds. Snow (1991) considers selected higher moments of the returns distribution. Bansal and Lehmann (1997) derive restrictions on $E[\ln(m)]$ that involve all higher moments of m and reduce to the HJ bounds if returns are lognormally distributed. Balduzzi and Kallal (1997) incorporate the implications for the risk premium on an economic variable. Cochrane and Hansen (1992) state restrictions in terms of the correlation between the stochastic discount factor and returns, while Cochrane and Saa'-Requejo (1996) bound the Sharpe ratios of assets' pricing errors. Hansen and Jagannathan

information can refine the HJ bounds, let $r_t \equiv R_t - R_{0,t} \mathbb{1}_e$ be the n -vector of excess returns, where R_0 is the riskless asset. We switch to excess returns in this section for intuition; in practice we use gross returns to avoid the loss of information implied by excess returns. For excess returns, Equation (1) says

$$E(m_t r_t | Z_{t-1}) = 0, \quad (4)$$

which is equivalent to:

$$E\{[m_t r_t] f(Z_{t-1})\} = 0 \text{ for all functions } f(\bullet), \quad (5)$$

where the unconditional expectation is assumed to exist. In other words, if we consider $r_t f(Z_{t-1})$ to represent the excess returns of “dynamic trading strategies,” then the presence of the conditioning information Z_{t-1} is essentially equivalent to the condition that $E(m_t r_t) = 0$ should hold, not only for the original excess returns r_t , but also for the dynamic trading strategies. The larger is the set of strategies for which the condition is required to hold, the smaller is the set of m_t 's that can satisfy the condition and the tighter are the bounds.

When there is no conditioning information, $f(\mathbb{1})$ is a constant and the SDF must price only the original excess returns. HJ observe that Equation (4) implies $E(m_t r_t \otimes Z_{t-1}) = 0$, choosing $f(\mathbb{1})$ to be the function $\mathbb{1} \otimes Z_{t-1}$. This “multiplicative” approach has become a standard in the asset pricing literature.

(1996) develop measures of distance between candidate SDFs and the m that would price the assets. Hansen, Heaton and Luttmer (1995) develop asymptotic distribution theory for specification errors on stochastic discount factors, where the HJ bounds are a special case.

HJ's multiplicative approach is not an efficient use of the lagged instruments. Ferson and Siegel (2000) derive *unconditionally-efficient* portfolios that have the maximum (unconditional) Sharpe ratio that can be obtained from a given set of returns, with portfolio weights that may depend on the conditioning information. Since the lower bound on SDFs is tighter when the maximum Sharpe ratio is higher, such portfolios represent an efficient use of the conditioning information. We use these results to construct *efficient portfolio* bounds, essentially choosing $f(\square)$ to be the set of all portfolio weight functions. The portfolio weight functions are restricted to sum to 1. Ferson and Siegel (2000) observe that the efficient portfolio weights are robust to extreme signals about asset returns, which suggests that estimates of the efficient portfolio bounds may have attractive sample properties.

Gallant, Hansen and Tauchen (GHT, 1990) provide a greatest lower bound from (1), which implies that (5) holds for *all* functions $f(\square)$. Bekaert and Liu (1999) show how to compute an optimal bound in a multiplicative framework, where the bound is shown to reach the GHT bound when the conditional moments are correctly specified. They point out that the GHT bounds are invalidated when the conditional moments involved in their computation are incorrectly specified.

In summary, the versions of HJ bounds that we study may be understood through Equation (5). First are the *multiplicative bounds* of Hansen and Jagannathan (1991), who choose $f(\bullet)$ to be an identity function. Second are the *efficient portfolio bounds*, where $f(\bullet)$ is a vector of portfolio weights that may depend on Z_t and sum to 1. Finally, the *optimal bounds* of Gallant, Hansen and Tauchen (1990) require Equation (5) to hold for all functions $f(\bullet)$. Note that we take the instruments as given; thus, we do not study how to choose Z_{t-1} .

3.1. Efficient Portfolio Bounds

Ferson and Siegel (2000) derive portfolios that use the given conditioning information, Z , to achieve unconditional mean-variance efficiency. These portfolios are used in our efficient-portfolio (UE) bounds. Let $x = x(Z) = (x_1(Z), \dots, x_n(Z))'$ denote the shares invested in each of the n assets, with the constraint that $x'e=1$. The observed gross return on the portfolio, $R_p = x'(Z)R$, has expectation and variance (using iterated expectations given Z to eliminate the unexpected returns) as follows:

$$\mu_p = E[x'(Z)\mu(Z)]$$

$$\sigma_p^2 = E\{x'(Z)[\mu(Z)\mu'(Z) + \Sigma_\varepsilon(Z)]x(Z)\} - \mu_p^2, \quad (6)$$

where $\mu(Z)$ denotes the conditional mean vector of the n returns, given Z , and $\Sigma_\varepsilon(Z)$ is the conditional covariance matrix. Define the following constants:

$$\alpha_1 = E\left(\frac{1}{e'\Lambda e}\right) \quad (7)$$

$$\alpha_2 = E\left(\frac{e'\Lambda\mu(Z)}{e'\Lambda e}\right) \quad (8)$$

$$\alpha_3 = E\left[\mu'(Z)\left(\Lambda - \frac{\Lambda e e'\Lambda}{e'\Lambda e}\right)\mu(Z)\right], \quad (9)$$

where

$$\Lambda = \Lambda(Z) = [E(RR'|Z)]^{-1} = [\mu(Z)\mu'(Z) + \Sigma_\varepsilon(Z)]^{-1}. \quad (10)$$

Proposition 2 (Ferson and Siegel 2000). Given unconditional expected return μ_P and n risky assets, the portfolio having minimum unconditional variance is determined by the optimal weights:

$$x'(Z) = \frac{e'\Lambda}{e'\Lambda e} + \frac{\mu_P - \alpha_2}{\alpha_3} \mu'(Z) \left(\Lambda - \frac{\Lambda e e'\Lambda}{e'\Lambda e} \right) \quad (11)$$

The variance of the portfolio defined by $x(Z)$ is

$$\sigma_P^2 = \left(\alpha_1 + \frac{\alpha_2^2}{\alpha_3} \right) - \frac{2\alpha_2}{\alpha_3} \mu_P + \frac{1 - \alpha_3}{\alpha_3} \mu_P^2 \quad (12)$$

The proof is given by Ferson and Siegel (2000).

To implement the UE bounds we must specify the conditional mean and variance functions, $\mu(Z)$ and $\Sigma_e(Z)$. The efficient set constants α_1 , α_2 , and α_3 are then estimated using sample averages. If these moments are incorrectly specified the portfolio weight given by (11) will no longer be efficient, but it still describes a valid dynamic portfolio strategy. The mean variance boundary and Sharpe ratio constructed from the strategy provides a valid but inefficient bound on stochastic discount factors in this case.

Fixed-weight combinations of any two portfolios on an unconditional mean-standard-deviation boundary can describe the entire boundary (Hansen and Richard, 1987). To form the UE bounds we pick two “arbitrary” portfolios. One is the global minimum-variance portfolio, which has the following mean and variance:

$$\mu^* = \alpha_2 / (1 - \alpha_3). \quad (13)$$

$$(\sigma^*)^2 = \alpha_1 - \alpha_2^2 / (1 - \alpha_3) \quad (14)$$

This follows by choosing μ_p to minimize the quadratic function for σ_p^2 , as given in (12). The second portfolio is chosen by setting μ_p equal to an arbitrary target return. For a given $\{\mu_p, \sigma_p\}$ hyperbola constructed from the two unconditionally efficient portfolio returns, the corresponding HJ bound can be obtained using Equation (3). Here, Σ is a 2×2 matrix and μ is a two-vector of the unconditional means of the two unconditionally efficient portfolios.

3.2. Optimal Bounds

This section provides a convenient, closed-form expression for the optimal HJ bounds that were originally derived by Gallant, Hansen and Tauchen (1990). First, define the following conditional-efficient-set constants, analogous to the efficient-set constants used in traditional mean-variance analysis (see, e.g., Ingersoll, 1987):

$$\begin{aligned}\alpha(Z) &= e' \Sigma_\varepsilon^{-1}(Z) e \\ \beta(Z) &= e' \Sigma_\varepsilon^{-1}(Z) \mu(Z) \\ \gamma(Z) &= \mu'(Z) \Sigma_\varepsilon^{-1}(Z) \mu(Z)\end{aligned}\tag{15}$$

Proposition 3: Optimal Hansen-Jagannathan Bounds (Gallant, Hansen and Tauchen, 1990). The stochastic discount factor m with minimum variance for its expectation $E(m)$ that satisfies $E(mR | Z) = e$ is given by

$$m = \zeta(Z) + [e - \zeta(Z)\mu(Z)]' \Sigma_\varepsilon^{-1}(Z) [R - \mu(Z)]\tag{16}$$

where $\zeta(Z)$, the conditional mean of m given Z , is defined as

$$\zeta(Z) = E(m|Z) = \frac{\beta(Z)}{1+\gamma(Z)} + \frac{1}{1+\gamma(Z)} \left\{ \frac{E(m) - E\left(\frac{\beta(Z)}{1+\gamma(Z)}\right)}{E\left(\frac{1}{1+\gamma(Z)}\right)} \right\} \quad (17)$$

and the unconditional variance of m is

$$\sigma_m^2 = \frac{\left[E(m) - E\left(\frac{\beta(Z)}{1+\gamma(Z)}\right) \right]^2}{E\left(\frac{1}{1+\gamma(Z)}\right)} + E[\alpha(Z)] - E\left[\frac{\beta^2(Z)}{1+\gamma(Z)}\right] - [E(m)]^2 \quad (18)$$

A proof of Proposition (3) is available by request to the authors. The result may be verified by computing $E(mR'|Z) = e'$ using Equation (16) for m , which holds for any definition of $\zeta(Z)$. Then, note that any other stochastic discount factor with the same unconditional mean as the m given by (16) can be expressed as $m + \varepsilon$ where $E(\varepsilon) = E(\varepsilon m) = 0$, and thus its variance is larger than the variance of m .

Equation (18) may be used directly to compute the optimal HJ bounds. As with the UE bound, it is necessary to specify the conditional mean function $\mu(Z)$ and the conditional variance function $\Sigma_\varepsilon(Z)$. The four unconditional expectations that appear in Equation (18) may be estimated from the corresponding sample means, independent of the value of $E(m)$. As emphasized by Bekaert and Liu (1999), if the moments are incorrectly specified the result may not be a valid bound on the variance of SDFs.

3.3. Discussion

The optimal bounds provide the greatest lower bound on stochastic discount factors. The UE bounds incorporate an additional restriction to functions of the conditioning information that are

portfolio weights, which sum to 1.0 at each date. This reduces the flexibility of the UE bounds to exploit the conditioning information, and thus they do not attain the greatest lower bound. Intuitively, suppose there was only one asset. Then the restricted weight could not respond at all to the conditioning information.

The additional restriction in the UE bounds may be understood in terms of the duality between Hansen-Jagannathan bounds and the usual mean-standard deviation diagram for returns. For a given value of $E(m)$, the value of σ_m on the HJ boundary is determined by the maximum squared Sharpe ratio when the implicit risk-free rate is $[E(m)]^{-1} - 1$. In the UE bounds the Sharpe ratio for a given $E(m)$ is achieved by a fixed-weight combination of the two unconditionally efficient portfolios, weighted according to the fixed value of $E(m)$. In the optimal bounds we choose $E(m|Z)$ for each realization of Z subject only to the limitation that $E[E(m|Z)]$ is the fixed $E(m)$. This allows the minimization to obtain the unrestricted optimal bound.³

While the UE bounds do not attain the greatest lower bound, they are nevertheless empirically interesting in view of two forms of “robustness.” The first, as emphasized by Bekaert

³ The difference between the two bounds may also be understood using the characterization of mean-variance frontiers from Hansen and Richard (1987). If a portfolio minimizes unconditional variance for a given mean, in the set of all returns that can be formed by trading with Z , it is unconditionally efficient (UE). Hansen and Richard show such UE portfolios are also conditionally mean-variance efficient (CE) for each realization of Z . Both the optimal and the UE bounds are formed from CE portfolios. Hansen and Richard show that any CE portfolio return is $w_t R_{1,t+1} + (1 - w_t) R_{2,t+1}$, where $R_{1,t+1}$ and $R_{2,t+1}$ are two UE returns which are also CE. For a given realization of Z_t , each point on the conditional mean-variance boundary implies a corresponding risk-free rate $[E(m|Z_t)]^{-1} - 1$ and a weight, w_t . In contrast, all UE portfolio returns can be formed as $w R_{1,t+1} + (1 - w) R_{2,t+1}$, where w does not depend on Z_t . In the UE bounds we fix w to correspond to $[E(m)]^{-1} - 1$.

and Liu (1999) is that the portfolio-based bounds remain valid when the conditional moments are incorrectly specified. Second, Ferson and Siegel (2000) show that the UE portfolio weights, unlike traditional mean-variance optimal weights, are “conservative,” in the sense that they avoid extreme positions in risky assets when the conditional moments are extreme. The UE portfolio weight, as a function of a conditional expected return, appears similar to the redescending influence curves used in robust statistics. Thus, the UE weights should be robust to outlier observations. These features may translate into robust sampling properties of the UE bounds.

4. Bias Correction

Consider first the case of no conditioning information, as in the fixed bounds. A simple bias-adjusted estimator assumes normally-distributed returns, and is based on standard results for exact finite sample distributions. Assume that T independent observations are made on the asset vector R . When the sample average $\hat{\mu}$ and the sample covariance matrix S (dividing by T) are used, we have the maximum likelihood estimate of the variance bound:

$$\hat{\sigma}_m^2 = [e - E(m)\hat{\mu}]' S^{-1} [e - E(m)\hat{\mu}]. \quad (19)$$

Assuming normality, the quadratic form in (19) has a noncentral chi-squared distribution, directly related to the distribution of a maximum squared Sharpe ratio, studied by Jobson and Korkie (1980). Using this distribution (also derived as a special case of Shanken 1982 and 1987, $k=0$) we find the mean of $\hat{\sigma}_m^2$. The estimated variance is biased upward (i.e., the true variance is overestimated). We solve for a transformation of $\hat{\sigma}_m^2$ that is unbiased.

Proposition 4. If asset returns are multivariate normal, then the expectation of the estimated variance of m in Equation (19) is given by

$$E(\hat{\sigma}_m^2) = \frac{n}{T-n-2} [E(m)]^2 + \frac{T}{T-n-2} \sigma_m^2 \quad (20)$$

and an unbiased estimator of the variance is given by

$$\hat{\sigma}_{m, adjusted}^2 = \left(1 - \frac{n+2}{T}\right) \hat{\sigma}_m^2 - \frac{n}{T} [E(m)]^2, \quad (21)$$

in the sense that $E(\hat{\sigma}_{m, adjusted}^2) = \sigma_m^2$.

This expression reveals the importance of the number of assets, n , relative to the number of time-series observations, T , for the determination of the bias. Approximately, the adjustment shrinks the estimated variance towards the value $-[E(m)]^2$, shrinking by the fraction n/T .

While the finite-sample adjustment in Equation (21) is developed for the case of no conditioning information, it may be applied to the multiplicative bounds of HJ. To see this, note that Equation (1) implies:

$$E(m_t R_t \otimes Z_{t-1} - e \otimes Z_{t-1} | Z_{t-1}) = 0. \quad (22)$$

Dividing the components of Z_{t-1} by their unconditional means and then taking the unconditional expectation implies:

$$E(m_t R_t \otimes \tilde{Z}_{t-1}) = e, \quad (23)$$

where $\tilde{Z}_{t-1} = Z_{t-1} / E(Z_{t-1})$ and $/$ denotes element-by-element division. Treating $R_t \otimes \tilde{Z}_{t-1}$ as the expanded set of “returns,” the multiplicative bounds are computed in the same fashion as the

fixed bounds. In the finite sample adjustment, n is taken to be the number of original assets times one plus the number of lagged instruments. The adjustment in this case is approximate, as it ignores the uncertainty due to the fact that $E(Z)$ must be estimated by the sample means. In our simulations we account for this uncertainty.

Building on Proposition 4, we provide approximate finite-sample bias corrections for the optimal and UE bounds.

Proposition 5. If asset returns are jointly normal, conditional on Z , and the maximum likelihood estimators for $E(R|Z)$ and $\Sigma_\varepsilon(Z)$ are used to form $\hat{\sigma}_{m^*}^2$ in the UE or optimal bounds, then an approximate bias-adjusted estimator is:

$$\hat{\sigma}_{m^*,adjusted}^2 = \left(\frac{T-n-2}{T} \right) \hat{\sigma}_{m^*}^2 - \frac{n}{T} [E(m)]^2 + \frac{2}{T} Var[E(m|Z)] \quad (24)$$

Proof: Both the UE and optimal bounds may be represented as the variance of a particular SDF, m^* , which may be expressed as

$$m^* = E(m|Z) + [e - E(m|Z)\mu(Z)]' \Sigma_\varepsilon(Z) [R - \mu(Z)]. \quad (25)$$

The optimal and UE bounds differ in the specification of the $E(m|Z)$ function. Computing the variance of (25),

$$\sigma^2(m^*) = Var[E(m|Z)] + E \left\{ [e - E(m|Z)\mu(Z)]' \Sigma_\varepsilon^{-1}(Z) [e - E(m|Z)\mu(Z)] \right\}. \quad (26)$$

For an estimated bound we replace $\mu(Z)$ and $\Sigma_\varepsilon^{-1}(Z)$ with their MLE estimates, which results in $\hat{\sigma}_{m^*}^2$. Assuming conditional joint normality of the returns, given Z , we evaluate the right-hand

term of Equation (26). Using iterated expectations, we first evaluate the conditional expectation given Z of the second term, taking $E(m|Z)$ as given and using Proposition 4. Then, taking the unconditional expectation of this result we arrive at the approximation

$$E(\hat{\sigma}_{m^*}^2) \cong \left(\frac{T}{T-n-2} \right) \sigma^2(m^*) + \left(\frac{n}{T-n-2} \right) [E(m)]^2 + \left(\frac{-2}{T-n-2} \right) Var[E(m|Z)]. \quad (27)$$

The approximation arises because we assume that the parameters in the $E(m|Z)$ function are at their probability limits in (27). Rearranging (27) as before we obtain the adjusted estimator.

Proposition 5 differs from Proposition 4 by the additional term, $(2/T)Var[E(m|Z)]$. Consistent estimates of this term for the optimal and UE bounds are obtained as the sample variances of the relevant expressions, evaluated at the MLE parameter estimates. For the optimal bounds, $E(m|Z)$ is specified in Equation (17). For the UE bounds, it can be shown that

$$Var[E(m_{UE}|Z)] = \left(\frac{1-E(m)\mu_p}{\sigma_p} \right)^2 \frac{Var[E(R_{UE}|Z)]}{\sigma_p^2}, \quad (28)$$

where R_{UE} is the portfolio formed with the weights given in (11), where μ_p and σ_p^2 are chosen to correspond to the portfolio located at the point on the UE portfolio frontier tangent to a line drawn from $[E(m)]^{-1}$ on the expected gross return axis.⁴

⁴ These values may be found by selecting μ_p to maximize $\{\mu_p - [E(m)]^{-1}\}^2 / \sigma_p^2$, where σ_p^2 is given by Equation (12).

5. Data

We use three different data sets in our empirical illustrations. An annual and a quarterly data set are constructed for comparability with Hansen and Jagannathan (1991). A monthly data set provides an example representative of more recent asset-pricing studies. Summary statistics are provided in Tables 1 and 2.

The annual data set used by Hansen and Jagannathan consists of real returns on a value-weighted stock index and short term real interest rates, from Shiller (1982). The annual data cover the 1891-1985 period. The lagged instruments consist of a constant and the first lagged values of the two real returns.⁵

The quarterly returns data also follow Hansen and Jagannathan, and they are the real, 3-month holding period returns on Treasury bills with initial maturities of 3, 6, 9, and 12 months; a total of four asset return series. The returns are computed from the yields reported in the Center for Research in Security Prices (CRSP) Fama files for original-issue twelve-month bills. Real returns are the nominal returns deflated by the component of the CPI relating to nondurable goods, as in Ferson and Harvey (1992). The quarterly data cover the period from the third quarter of 1964 through the fourth quarter of 1987, which is the same as HJ. The lagged instruments, following HJ, consist of a constant and the first lagged values of the real returns and real, per capita consumption growth, which we obtain from the Commerce Department via Citibase.

Our monthly data set includes the total returns (price change plus dividends) on twenty five industry portfolios from Harvey and Kirby (1996), measured for the period February, 1963 to

⁵ These data are published in Shiller (1989), Chapter 26, Tables 26.1-26.2.

December, 1994.⁶ The portfolios are created by grouping individual common stocks according to their SIC codes and forming value-weighted averages (based on beginning-of-month values) of the total returns within each group of firms. Table 2 shows the industry classifications for the 25 portfolios and their summary statistics. The instruments are: (1) the lagged value of a one-month Treasury bill yield, (2) the lagged dividend yield of the Standard and Poors 500 (S&P500) index.

6. Empirical Results

In this section we present sample versions of the various HJ bounds. The “fixed” bounds use no conditioning information, and so are determined by a fixed-weight combination of the basic asset returns, as in Equation (2). We compute maximum likelihood estimates of the bounds using the sample mean vector and covariance matrix. Under more general assumptions these are consistent method-of-moment estimates.

To form the efficient portfolio and optimal bounds we must specify the conditional mean function $\mu(Z)$ and the conditional variance function $\Sigma_\varepsilon(Z)$. Initially, we simply regress the returns on a constant and the lagged instruments. The fitted values of the regression are taken as $\mu(Z)$ and the sample covariance matrix of the residuals is our estimate of the (fixed) conditional covariance matrix. We estimate the portfolio constants α_1 , α_2 , and α_3 , replacing the unconditional expectations by the corresponding sample averages. For the optimal bounds, we use Equation (18) to define the minimum variance boundary, where the unconditional expectations are estimated by their corresponding sample means. For the UE bounds, we use

⁶ We are grateful to Campbell Harvey for providing the data.

Equation (11) to form two portfolios. One sets the target mean μ_p equal to the grand mean of the asset returns, the other is the global minimum variance portfolio, with target mean given by Equation (13).

6.1. Economic Significance

Figure 1 illustrates the results of estimating the various bounds using the monthly data set over the July 1963 - December 1994 period with no finite sample adjustments. A valid stochastic discount factor must lie above the bounds, “in the cup.” The bounds using conditioning information plot above the fixed bounds, illustrating that conditioning information allows one to obtain a seemingly-more-powerful diagnostic, ruling out more stochastic discount factors. Also, there are substantial differences between the various bounds, which suggests that the choice of which bound to use is a substantive decision.




Figure 1. Unadjusted HJ bounds from monthly data, July 1963 - December 1994 with the estimated conditional model (“FFC”) and unconditional model (“FFU”).

A striking feature of Figure 1 is that the sample multiplicative bound plots *above* the sample optimal bound. The theory of Section 3 implies that this cannot occur, except as a result of misspecification or finite sample error. Thus, the figure motivates a study of the sampling properties of the bounds.

Figure 2 applies the finite sample adjustments to the bounds. Now the ordering of the various bounds appears reasonable, with the optimal bounds plotting above. The effect of the finite-sample bias adjustment to the multiplicative bounds appears substantial. This suggests that there may be substantial finite-sample biases.

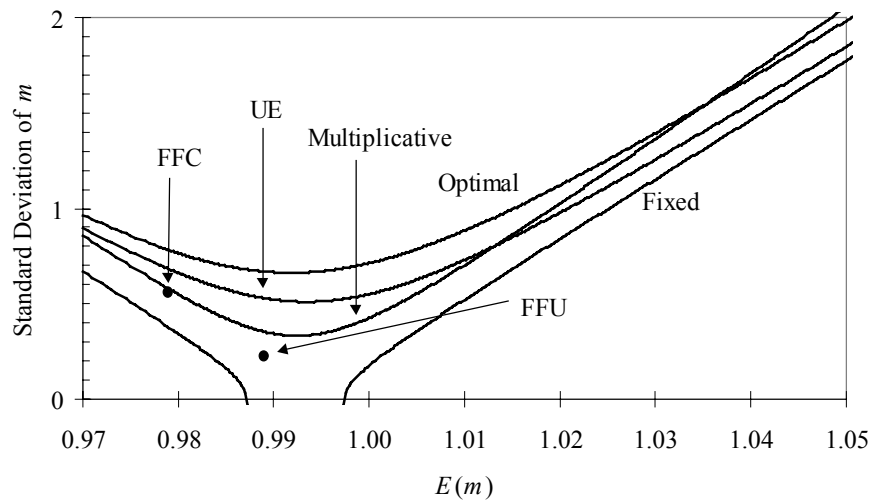


Figure 2. Bias-adjusted HJ bounds from monthly data, July 1963 - December 1994 with the estimated conditional model (“FFC”) and unconditional model (“FFU”).

The differences between the various bounds, and the effects of the finite-sample adjustment appear substantial, but it is important to address the economic significance of these effects. Using an example from the Lucas (1978) consumption-based model, we find that the consumption SDF is

so far outside the bounds that our results would not change inferences about the validity of that model. A simple model of habit persistence, like the one used by HJ, produces a similar result. However, when the SDFs of interest are closer to the bounds the results will be economically significant. There are reasons to think that many interesting SDFs are close to the bounds. Intuitively, the boundary of the admissible region for SDFs is formed from a combination of asset returns. Many SDFs in the literature are also a combination of asset returns. To illustrate, we consider the example of the “three-factor model” of Fama and French (1993, 1996).

Fama and French advocate a linear model in which three return factors describe the SDF. The factors are a market portfolio return, the difference between the returns of a small-stock and a large-stock portfolio (SMB) and the difference between a high and a low book-to-market portfolio (HML). While there is some controversy over the justification for this model it has been popular in recent studies.

A linear beta pricing model implies an SDF in which m_t in Equation (1) is a linear function of the factors (see Dybvig and Ingersoll, 1982, or Ferson and Jagannathan, 1996): $m_t = a(z) + B(z)' F_t$, where F_t is the vector of factors. Farnsworth *et al* (2000) estimate SDF formulations of the Fama-French model using a monthly data set for July, 1963 - December 1994, only five months shorter than our sample period, and we use their results here. Following Cochrane (1996) they assume that the coefficients $a(z)$ and $B(z)$ are linear functions of the lagged instruments, a Treasury bill and a dividend yield. In this case we have a *conditional* version of the model, which we denote by “FFC” in the graphs. When $a(z)$ and $B(z)$ are constants, so that no lagged instruments are used in forming the SDF, we have an *unconditional* model (“FFU”).

In Figure 1, with no bias adjustments, the Fama-French SDFs plot close to the fixed bounds but below the bounds with conditioning information. While no standard errors are shown, if we superimpose the sampling variation from the subsequent tables, it is likely that the model would be rejected using the biased bounds with conditioning information.

In Figure 2, with the bias adjustments applied, the SDFs are the same as in the previous figure.⁷ The SDFs now plot close to the multiplicative bounds. After bias adjustment one would not reject the Fama-French SDF using the multiplicative bounds, as the SDFs are within two standard errors of the adjusted bounds. This reverses the conclusion from Figure 1, illustrating that the biases in the bounds are economically significant. Unlike the multiplicative bounds, Figure 2 suggests that the optimal bounds would be likely to reject the Fama-French SDF. The point FFC lies 2.02 standard errors below the optimal bound, while FFU is 4.15 standard errors below.⁸ Thus, the choice between bounds with conditioning information is also a matter of economic significance.

⁷ For some values of $E(m)$ the adjusted $\hat{\sigma}_m^2$ can be less than zero, and this occurs in Figure 2 for the fixed bounds. As we describe below, the fixed bounds have little economic content when sampling variation is accounted for, which confirms the importance of conditioning information. The local concavity of the fixed-bound curve for values of $\hat{\sigma}_m$ near zero is explained by the fact that the square root function has infinite slope at zero.

⁸ These examples ignore the sampling error in the location of the FFC and FFU points, and should therefore be only taken as illustrative. See Burnside (1994) and Tierens (1993) for examples that account for only this source of sampling error in specific SDF models.

7. Simulation Results

We conduct a simulation study of the sampling properties of the various bounds. The experiments accommodate data that are dependent over time. For example, the monthly dividend yields and lagged interest rates have first order sample autocorrelations in excess of 0.9, while asset returns have smaller autocorrelations as may be seen in Tables 1 and 2. We focus on capturing the autocorrelation of the lagged instruments. Using these instruments to model the conditional means of the asset returns, the simulated asset returns inherit mild serial dependence.

We estimate a first-order VAR for the lagged instruments, and we use the estimated coefficient matrix as parameters of the model. The parameters for the conditional means of the asset returns are estimated by regressions on the lagged instruments. The residuals from the VAR, U_Z , and the deviations of the asset returns from their conditional means, U_r , represent the shocks in the model. We concatenate these as (U_r, U_Z) and compute the sample covariance matrix as a parameter of the simulation. We generate the artificial shocks in the simulations by drawing data with this covariance matrix, either from a normal distribution or by resampling from the actual residuals. We build up the time series of the simulated instruments recursively using the VAR coefficients and the simulated U_Z shocks. We find that the artificial instruments have autocorrelations very similar to the sample data. The artificial returns are formed as the conditional means, evaluated at the values of the artificial instruments, plus an independent draw from the U_r distribution. (In a later section we explore the robustness of the results to more general data generating schemes.)

In each simulation trial the bounds are estimated from a sample of artificial data, in the same way that we estimate the bounds in the previous section. The estimated HJ curve may be described in any given example by the values of the three coefficients, a , b , and c , as

$$\hat{\sigma}^2(m) = a + b[E(m) - 1] + c[E(m) - 1]^2. \quad (29)$$

Thus, for each simulation trial we record the values of a , b , and c . We use 5,000 Monte Carlo trials for each data set. For a given $E(m)$ the values of a , b , and c , that summarize a simulation trial, determine a value of $\hat{\sigma}_m^2$. The 5,000 values of $\hat{\sigma}_m^2$, one for each simulation trial, are used to produce the summary statistics shown in the tables. The number of observations in each of the artificial samples is equal to that of an actual data set. For example, in the quarterly data set each of the 5,000 trials uses an artificial sample of 93 quarters, representing the four assets and instruments. We conduct simulations corresponding to the annual, quarterly and monthly data sets. Subsequent tables summarize the results of the simulations at $E(m)=1$. (Graphical summaries showing a range of $E(m)$ values are available by request to the authors).

7.1. A Benchmark: The “True” HJ Bounds

Since we do not wish to tie our results to a particular model economy (and corresponding SDF) we use large-scale simulation to find the benchmarks against which the sampling properties are measured. In order to determine the “true” bounds, we form artificial samples just like in the simulations, but with one million observations. Values of the true bounds are the sample values in the artificial sample with 1,000,000 observations. (Averaging across 100 simulations with 10,000 observations produces similar results.)

The true variance bounds are shown in Table 3, where the simulations use normally distributed data. A comparison across the three panels shows that the bounds are the highest in the quarterly data set. This reflects the high Sharpe ratios that appear in samples of Treasury bill returns, consistent with HJ (1991).

Comparing the different bounds for the same data set confirms that, abstracting from sampling error, the different bounds can produce vastly different results. For example, in monthly data the multiplicative bound for σ_m^2 is about three times the fixed bound. The optimal bound is almost four times the fixed bound, suggesting that the way in which conditioning information is used also makes a difference. The differences are also large in the quarterly data set. In the annual data, the differences between the various bounds are relatively small.

7.2. The Location of the Sample Bounds

Table 3 reports the mean of the estimated bounds, $\hat{\sigma}_m^2$, taken across the 5,000 simulation trials. Comparing these values with the true bounds shows the expected finite sample bias. All of the bounds display an upward bias; i.e., the sample bounds are higher than the true bounds. Thus, some valid stochastic discount factors are expected to plot outside of the sample HJ bounds. The upward bias of the fixed bounds is consistent with the previous studies of Burnside (1994) and Tierens (1993). Table 3 extends the evidence to the bounds with conditioning information. The expected sample bounds range from 113% to 200% of the true bounds on σ_m^2 .

Consistent with Propositions 4 and 5, the finite sample biases are more extreme where the number of time series observations is small relative to the number of assets. For example the annual data include 95 observations on 2 assets, and the ratio of estimated to true bounds is 113% to 129%. The monthly data include 383 observations on 25 assets, and the ratio of estimated to true bounds is

159% to 200%. The quarterly data presents an intermediate case. In each data set the multiplicative bounds have the largest bias; in the quarterly and monthly data the differences are substantial. The ratio of the estimated to the true bound is slightly closer to 1.0 for the UE than for the optimal bound.

The relation of the Hansen-Jagannathan bounds to the maximum squared Sharpe ratio provides intuition for the extreme sampling bias of the multiplicative bounds, compared with the other bounds. It is well known from the classical mean-variance analysis that portfolios based on the usual MLE estimates of the mean returns and their covariance, tend to be biased in favor of overstated Sharpe ratios. This intuition is reflected in Propositions 4 and 5. The bias is greater, for a given sample size T , when more assets are included in the portfolio.⁹ Since the multiplicative bounds create additional “assets,” the maximum Sharpe ratio and, thus, the bound on σ_m^2 , is likely to be more upwardly biased. The smaller finite-sample bias of the UE bound, in contrast, reflects the robustness of UE portfolios, discussed by Ferson and Siegel (2000).

7.3. Bias Adjustment

Table 3 reports the averages of the bounds, adjusted for finite sample bias, taken across the 5,000 simulation trials (Adj. Mean). A comparison with the true bounds and the unadjusted means shows the effectiveness of the adjustment to the location of the bounds. The expected adjusted bounds range from 90% to 131% of the true bounds. The range is tighter in the annual data, where

⁹ See Frost and Savarino (1988), MacKinlay (1987), and Green and Hollifield (1992) for illustrations.

all adjusted bounds are within 10%, and the widest in the monthly data, where the adjustment factor is larger.

The adjustments provide a dramatic improvement in quarterly and monthly data, especially for the fixed and multiplicative bounds. For example the adjustment reduces the quarterly multiplicative bound bias from 71% to 8.5%, and the monthly multiplicative bound bias from 100% to -3%. It performs even better on the fixed bounds. Since we draw normally distributed data in these simulations, the adjustment to the fixed bound would be exact if the data were serially independent. In the multiplicative bound the products of normals imply nonnormal dynamic strategy returns, which reduces the accuracy of the adjustment. (We consider nonnormal and heteroskedastic artificial data below.)

The adjustments for finite sample bias also provide significant improvements in the location of the optimal and UE bounds. Before adjustment, the biases of these bounds range from 28% to 59% in the quarterly and monthly data. After adjustment the biases are roughly halved.

7.4. The Precision of the Bounds

The value of the bounds as a diagnostic tool also depends not only on their location, but also on their precision. The simulations provide information on the sampling variation of the bounds. Table 3 reports the standard deviations of the bounds, evaluated at $E(m)=1$, with and without adjustment, taken across the 5,000 simulation trials. Recall that the adjustment approximately shrinks the uncentered second moment of the stochastic discount factor, multiplying it by a factor less than 1.0. This results in smaller standard errors. Thus, our finite sample adjustments provide another benefit, in addition to reducing the bias in the location of the bounds. They also produce bounds with greater precision.

Comparing the various bounds for a given data set, the fixed bounds have the smallest standard deviations. Of course, they also have the smallest average values. Among the bounds with conditioning information, the UE bounds always have the smallest standard deviations. Thus one appeal of the UE bounds is their relative precision.

Table 3 shows that none of the bounds have much economic content in the annual data set. Negative values of σ_m^2 lie within two standard deviations of the true version of each bound. In the quarterly and annual data sets the bounds do place substantive restrictions on σ_m^2 . In these data sets the bounds with conditioning information are much more restrictive of SDF variances than the fixed bounds. In particular, the fixed bound in monthly data has virtually no economic content, as the true bound is within two standard errors of zero. This result was also illustrated in Figure 2.

The efficient portfolio and optimal bounds are more restrictive of SDF variances than the multiplicative bounds. For example, in the monthly data set a value of $\sigma_m^2 = 0.128$ is two standard deviations below the true multiplicative bound. For the efficient portfolio bound the corresponding value is 0.148, while for the optimal bound it is 0.173. In the quarterly sample the UE and optimal bounds have even larger advantages. Thus, when we consider both the location and the precision of the bounds, the efficient use of the conditioning information produces markedly tighter bounds, relative to the standard multiplicative approach. This reinforces the impression from Figure 2 that it is important to efficiently use the conditioning information in variance bounds.

8. Effects of Nonnormality and Heteroskedasticity

The previous simulations use strong assumptions about the data generating process. These include (1) normality and, consistent with normality, (2) homoskedasticity. There is evidence in the literature inconsistent with each of these assumptions. In this section we explore the robustness of the results to the use of alternative assumptions.

We conduct two experiments in which we generalize the data generating process by progressively relaxing the assumptions (1) and (2). This provides information on how robust our previous results are to alternative data generating schemes. For each experiment we conduct a new large scale simulation to define a “true” bound, as the true bound may depend on the specification of the moments in the data generating process.

8.1. Nonnormality

In the first experiment we relax the assumption that the shocks in the data generating process are normally distributed. Instead of drawing normally distributed shocks with a given covariance matrix, we use an approach similar to the bootstrap (see, e.g. Efron, 1982). We resample vectors from the sample of residuals (U_r, U_z) , choosing dates randomly with replacement. The artificial data are otherwise generated as before. The simulated data will be homoskedastic but not normally distributed, on the assumption that the sample is not normally distributed.

The results of the first experiment are summarized in Panels A-C of Table 4. The true values of the fixed-weight bounds are the same as in Table 3, because the fixed bound is a consistent estimator and no lagged instruments are used. The true bounds that use lagged instruments are affected only very slightly by nonnormality. The other results are also similar to those in Table 3.

Among the bounds with conditioning information, the multiplicative bounds have the largest bias, and the UE bounds the smallest. The UE bounds have the smallest standard errors. Using the true location and accounting for the standard deviations, the UE and optimal bounds are more restrictive of the data than the multiplicative bounds. The performance of the finite sample adjustment is consistent with our previous observations. Its performance is not degraded by the nonnormality in the data.

8.2. Heteroskedasticity

Heteroskedastic data raises some new issues. First, the expected values of objects like $\mu'(Z)\Sigma_\varepsilon^{-1}(Z)$, which appear in both the optimal and UE bounds, will differ from their expectations under homoskedasticity if the conditional mean and elements of the conditional covariance matrix are correlated. Thus, the “true” locations of the bounds are expected to shift. Second, the issue of correctly specifying the heteroskedasticity becomes potentially important. In our previous simulations the artificial econometrician essentially uses the correct data generating process, but needs to estimate the parameters. Under heteroskedasticity, the correct process may not be obvious. If the wrong specification is used, the estimated optimal bounds may not be valid, and the UE bounds will be inefficient. In our simulations we assume that the artificial econometrician has correctly identified the data generating process to be the one we actually use. However, the parameters of the process must be estimated on the artificial data. In this sense, the experiment is symmetric with the previous cases. (See Bekaert and Liu, 1999, for cases where the data generating process may be incorrectly specified.)

In the second experiment we allow for both nonnormality and conditional heteroskedasticity in the data generating process. There are many ways to model conditional heteroskedasticity, and

we experimented with several alternatives. Given the large size of the conditional covariance matrix (25×25 in the monthly data) relative to the number of time-series observations (383) standard approaches such as ARCH and GARCH seem impractical.

We adopt a constant correlation structure for the conditional covariance matrix. The correlations are taken from the sample correlation matrix of the regression residuals U_r , formed as in the first two experiments. The conditional standard deviations are modeled by regressing the absolute residuals on the lagged instruments and saving the coefficients as parameters. For a given draw of Z in the simulation the conditional standard deviation is the regression coefficient applied to that Z , multiplied by $\sqrt{\pi/2}$. This approach is advocated by Davidian and Carroll (1987) and is similar to the approaches of Ferson and Foerster (1994) and Schwert and Seguin (1990). The conditional covariance is then formed as the constant correlation, multiplied by the conditional standard deviations implied by the value of Z . Modeling the conditional heteroskedasticity as described above, the generated data do not match the moments of the sample data as closely as in the previous experiments. We shrink the estimated time-varying covariance matrix towards the fixed covariance matrix to obtain a better match.¹⁰

Panels D-F of Table 4 show the results of the second experiment. Compared to the previous experiments the locations of the true bounds are raised slightly in most of the cases. The ratios of

¹⁰ The regression model for the heteroskedasticity produces returns whose regression R^2 on Z are implausibly high. We therefore shrink the conditional variance matrix towards the unconditional covariance matrix of the returns. In the annual and monthly data we use the simple average of the two. This produces artificial returns that more closely match sample regression results. In the quarterly data we find that a shrinkage factor of 0.5 is too small, and a shrinkage factor of 0.9 toward the fixed covariance of U_r is used.

the estimated to the true bounds, however, are similar. The biases vary from 15% to 107% across the cases, and the relative performance of the different bounds is similar to that observed before. The multiplicative bounds have the largest bias. The UE bounds have the smallest standard deviations, among the bounds with conditioning information. The finite-sample bias adjustments continue to perform well. The adjusted location of each estimated bound is within 11% of the true bound in the annual and quarterly data, except in the case of the optimal bound in quarterly data (18%). Even in monthly data the bias is cut in half, or better, by the adjustment. Accounting for sampling variation, the optimal and UE bounds are more restrictive of SDF variances than the multiplicative bound, confirming the importance of the efficient use of conditioning information in the bounds.

9. Asymptotic Standard Errors

Hansen, Heaton and Luttmer (HHL, 1995) derive an asymptotic standard error for the minimum second moment of a stochastic discount factor, in the case where there is no conditioning information. In the appendix we modify their results to use asset returns instead of asset prices and payoffs, and to accommodate arbitrary values of $E(m)$. A consistent estimator for the asymptotic variance of $\hat{\sigma}_m^2$ is obtained as the asymptotic variance of $(1/T) \sum \phi_t$, where

$$\phi_t = \left\{ -\left[\hat{\alpha}'(R_t - \hat{\mu}) \right]^2 - 2\hat{\alpha}' \left[E(m)R_t - e \right] \right\}, \quad (30)$$

where $\hat{\alpha} = S^{-1} \left[e - E(m)\hat{\mu} \right]$. We estimate the asymptotic variance of $(1/T) \sum \phi_t$ using the Newey-West (1987) estimator. The number of autocovariance terms is determined by examining

the sample autocorrelations of ϕ_t , and including the lags where the sample autocorrelations exceed two approximate standard errors.¹¹

Table 5 summarizes our evaluation of the sampling properties of the HHL asymptotic standard errors for the fixed variance bounds.¹² The “Empirical SD” and “Adjusted Empirical” are the standard deviations from the simulations, repeated from the previous tables for convenience. “Average Asymptotic” is the mean value of the HHL standard errors, taken over the 5,000 trials and “Adjusted Asymptotic” refers to the asymptotic standard error of the adjusted bounds. The table shows that the asymptotic standard errors are fairly reliable in the annual and quarterly data. They are understated, but by less than 10%, in each of the three experiments. In the monthly data, where the number of time series (383) is small relative to the number of assets (25), the asymptotic standard errors are less reliable. They are understated by almost 20% when the data are homoskedastic, and by as much as 44% under heteroskedasticity. In the monthly data the levels of the bounds are smaller numbers, on the order of 0.05 (versus 0.10 in annual and quarterly data), so a given absolute error represents a larger percentage bias.

¹¹ This criterion results in 12 lags in the monthly data, and zero lags in the annual and quarterly data.

¹² We also studied extensions of HHL’s results to the multiplicative and optimized bounds, obtaining an approximate solution in the multiplicative case. Simulations showed that first-order approximations are not sufficiently accurate to obtain a reliable standard error.

10. Summary and Conclusions

This paper offers a number of refinements and insights into the volatility bounds on stochastic discount factors, first developed by Hansen and Jagannathan (HJ, 1991). When there is no conditioning information the bounds are formed from fixed-weight combinations of the asset returns, with weights depending on the sample mean and covariance matrix. In the presence of conditioning information most studies in the literature have followed Hansen and Jagannathan, multiplying the returns by the lagged variables and constructing bounds with these “dynamic strategy” returns. We find that sample values of these bounds are upwardly biased, the bias becoming substantial when the number of assets is large relative to the number of time-series observations. This means that studies using the biased bounds run a risk of rejecting too many models for the stochastic discount factor. We argue that the magnitude of the bias is economically significant. We provide a finite-sample adjustment for this bias.

We compare Hansen and Jagannathan’s “multiplicative” approach with two alternative approaches to the use of conditioning information. One approach, following Ferson and Siegel (2000), is based on unconditionally mean-variance efficient dynamic trading strategies. We call these the “efficient-portfolio” (UE) bounds. The second approach, based on Gallant, Hansen and Tauchen (1990), provides the theoretically tightest possible bounds. We present a closed-form solution for this “optimal” bound, which simplifies the implementation and analysis. We also evaluate asymptotic standard errors for the HJ bounds, derived by Hansen, Heaton and Luttmer (1995). Our simulation study leads to several conclusions.

(1) *Multiplicative bounds* are easy to use but they can be terribly biased. Our finite-sample adjustment improves their specification, in the sense that the expected bias in the location of the

adjusted bounds is small. However, the sampling variation in the multiplicative bounds is larger than in the other bounds with conditioning information, and the bounds are less restrictive of SDF variances once sampling error is taken into account. If we could only use one version of the bounds with conditioning information, we would eliminate the multiplicative bound, based on these results.

(2) *Optimal bounds* are more difficult to implement than the multiplicative bounds, requiring a specification for the conditional means and variances of the asset returns. The magnitudes of the finite-sample biases are less than in the multiplicative case. The bias adjustment is the least effective for these bounds. However, accounting for sampling error the optimal bounds are the most restrictive of SDF variances. Based on these results we would prefer to use the optimal bound in a setting where we had a high degree of confidence in the specification of the conditional means and variances of the asset returns.

(3) *Efficient-portfolio bounds* are similar in complexity to the optimal bounds, also requiring a specification of the conditional moments. However, unlike the optimal bounds they are theoretically robust to an incorrect specification of the conditional moments (Bekaert and Liu, 1999). The UE bounds have smaller standard errors than either the multiplicative or the optimal bounds. Using our adjustments, the finite-sample bias is smaller than in the case of the optimal bounds. However, the UE bounds are not as restrictive of SDF variances as the optimal bounds. Based on these results, we advocate the UE bounds in settings where robustness and precision of the bounds are the important concerns.

(4) *Asymptotic Standard Errors* for the fixed bounds are mildly understated in our annual and quarterly samples, by less than 10%. The degree of bias is worse when the number of assets

is large relative to the number of time series observations. In our monthly sample, where the level of the bounds is a smaller number, the percentage bias is 20% or larger.

Appendix

Asymptotic Standard Errors: Fixed Bounds

This shows how we modify results from HHL (1995). They work in terms of asset prices and payoffs, while we use gross asset returns, which are more likely to be stationary. Following HHL (1995) consider the following two optimization problems.

$$\underset{m}{\text{Min}} \text{Var}(m) \text{ such that } E(mR) = 1, E(m) = m_0 \quad (\text{A.1})$$

$$\sigma_m^2 \equiv \underset{\alpha}{\text{Max}} \left\{ -E \left([\alpha'(R - \mu)]^2 \right) - 2\alpha'(m_0\mu - e) \right\}. \quad (\text{A.2})$$

The optimized values of the two problems are identical. To see this, recall that the stochastic discount factor that solves (A.1) is of the form $m = m_0 + \alpha^*(R - \mu)$ where $\alpha^* = \Sigma^{-1}(e - m_0\mu)$, as in Equation (2). From the first-order conditions to problem (A.2), the optimal value of α is the same α^* , and at this value

$$\sigma_m^2 = (e - m_0\mu)' \Sigma^{-1} (e - m_0\mu), \quad (\text{A.3})$$

which is also the solution to (A.1), as in Equation (3).

To obtain the asymptotic variance for the sample version of the bounds, define

$$\phi_t(\alpha) \equiv \left\{ -[\alpha'(R_t - \mu)]^2 - 2\alpha'(m_0R_t - e) \right\} \quad (\text{A.4})$$

Similarly let $\hat{\phi}_t(\alpha)$ be the value of (A.4) when the sample mean, $\hat{\mu}$, replaces μ , and let $\hat{\alpha} \equiv S^{-1}(e - m_0\hat{\mu})$ be the optimum value of α , evaluated at the MLE estimates $(S, \hat{\mu})$. We may then express the sample estimate of the fixed bound as $\hat{\sigma}_m^2 = (1/T) \sum_t \hat{\phi}_t(\hat{\alpha})$, and the true value as $\sigma_m^2 = E[\phi_t(\alpha^*)]$. We use the following decomposition:

$$\sqrt{T}(\hat{\sigma}_m^2 - \sigma_m^2) = \sqrt{T} \left\{ \left[\frac{1}{T} \sum_t \hat{\phi}_t(\hat{\alpha}) - \frac{1}{T} \sum_t \hat{\phi}_t(\alpha^*) \right] + \left[\frac{1}{T} \sum_t \hat{\phi}_t(\alpha^*) - E[\phi_t(\alpha^*)] \right] \right\} \quad (\text{A.5})$$

The first term of (A.5) isolates the uncertainty due to the estimate of the optimal α , while the second captures the time-series variation in $\hat{\phi}_t$, given the population value, α^* .

The first term of (A.5) converges in probability to zero under standard assumptions, and therefore does not contribute to the asymptotic variance of the estimator. To see this, consider an exact, second-order Taylor series expansion of $(1/T) \sum_t \hat{\phi}_t(\alpha^*)$ about $(1/T) \sum_t \hat{\phi}_t(\hat{\alpha})$. Since $\hat{\alpha}$ is optimal at the sample estimates, the first-order condition for problem (A.2) implies that the first term in the Taylor expansion is zero. Using the second term of the expansion and the first-order condition for problem (A.2), we have:

$$\sqrt{T} \left\{ \frac{1}{T} \sum_t \hat{\phi}_t(\hat{\alpha}) - \frac{1}{T} \sum_t \hat{\phi}_t(\alpha^*) \right\} = \frac{\sqrt{T}}{2} \left\{ (e - m_0\hat{\mu})' - (e - m_0\mu)' \Sigma^{-1} S \right\} (\hat{\alpha} - \alpha^*) \quad (\text{A.6})$$

The right-hand side of (A.6) is the product of two terms. The first converges in distribution to a normal with mean zero under standard assumptions. The second term, $(\hat{\alpha} - \alpha^*)$, converges almost surely to zero. Thus, the product of the two terms converges in probability to zero.

The preceding arguments imply that the asymptotic variance of $\sqrt{T}(\hat{\sigma}_m^2 - \sigma_m^2)$ is the asymptotic variance of $\sqrt{T}\left[(1/T)\sum_t \hat{\phi}_t(\alpha^*)\right]$. We obtain consistent estimates of the asymptotic variance by replacing α^* with a consistent estimator, $\hat{\alpha}$, as described in the text, and using the time series of the $\hat{\phi}_t(\hat{\alpha})$ to estimate the spectral density at frequency zero.

References

- Bansal, Ravi and Bruce N. Lehmann, 1997, Growth-Optimal Restrictions on Asset Pricing Models, *Macroeconomic Dynamics* **1**, 1-22.
- Bekaert, Geert, and Jun Liu, 1999, Conditioning Information and Variance Bounds on Pricing Kernels, working paper, Stanford University.
- Balduzzi, Pierluigi, and Heidi Kallal, 1997, Risk Premia and Variance Bounds, *The Journal of Finance* **52**(5), 1913-49.
- Burnside, Craig, 1994, Hansen-Jagannathan Bounds as Classical Tests of Asset-Pricing Models, *Journal of Business and Economic Statistics*, vol. **12**, no. 1, 57-79.
- Cochrane, John H., 1996, A Cross-Sectional Test of a Production-Based Asset Pricing Model, *Journal of Political Economy* **104**(3), 572-621.
- Cochrane, John H., and Lars P. Hansen, 1992, Asset Pricing Explorations for Macroeconomists, in O. J. Blanchard and S. Fischer, *NBER Macroeconomic Annual 1992*, **7**, 115-165.
- Cochrane, John H., and Jesus Saa'-Requejo, 1996, Beyond Arbitrage: Good-Deal Pricing of Derivatives in Incomplete Markets, working paper, University of Chicago.

- Davidian, Marie and R. J. Carroll, 1987, Variance Function Estimation, *Journal of the American Statistical Association* **82**, 1079-91.
- Dybvig, P. and J. Ingersoll, 1982, Mean variance theory in complete markets, *Journal of Business* **55**, 233-51.
- Efron, Bradford, 1982, *The Jackknife, the Bootstrap, and Other Resampling Plans*, Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Epstein, L., and S. Zin, 1991, Substitution, risk aversion and the temporal behavior of asset returns, *Journal of Political Economy* **99**, 263-286.
- Fama, Eugene F. and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**, 3-56.
- Fama, Eugene F. and Kenneth R. French, 1996, Multifactor Explanations of Asset Pricing Anomalies, *Journal of Finance* **51**, 55-87.
- Farnsworth, H., W. Ferson, D. Jackson and S. Todd, 2000, Performance Evaluation with Stochastic Discount Factors, working paper, University of Washington.
- Ferson, Wayne E., 1995, Theory and Empirical Testing of Asset Pricing Models. *Handbooks in OR & MS, Vol. 9, Chapter 5*. Amsterdam: Elsevier.
- Ferson, Wayne E. and Stephen R. Foerster, 1994, Finite Sample Properties of the Generalized Methods of Moments Tests of Conditional Asset Pricing Models, *Journal of Financial Economics* **36**, 29-56.
- Ferson, Wayne E., and Campbell R. Harvey, 1992, Seasonality and Consumption-Based Asset Pricing Models, *Journal of Finance* **47**(2), 511-552.

Ferson, Wayne E. and R. Jagannathan, 1996, Econometric Evaluation of Asset Pricing Models, Chapter 1 (pp.1-30) in the *Handbook of Statistics: vol. 14: Statistical Methods in Finance*, G.S. Maddala and C.R. Rao (editors), North Holland ISBBN: 0-444-81964-9.

Ferson, Wayne E., and Andrew F. Siegel, 2000, The Efficient Use of Conditioning Information in Portfolios, forthcoming, *Journal of Finance*.

Frost, P. and J. Savarino, 1988, For Better Performance: Constrain Portfolio Weights, *Journal of Portfolio Management* **15**(1), 29-34.

Gallant, A. Ronald, Lars P. Hansen and George Tauchen, 1990, Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics* **45**, 141-179.

Green, R., and B. Hollifield, 1992, When will mean variance efficient portfolios be well diversified? *Journal of Finance* **47**, 1785-1809.

Hansen, Heaton and Luttmer, 1995, Econometric Evaluation of Asset Pricing Models, *Review of Financial Studies* **8**(2), 237-274.

Hansen, L. P., and R. Jagannathan (1991). Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* **99**, 225-262.

Hansen, L. P., and S. F. Richard, 1987, The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica* **55**, No. 3, 587-613.

Harvey, Campbell R. and Chris Kirby, 1996, Analytic tests of factor pricing models, Unpublished working paper, Duke University, Durham, N.C.

- Ingersoll, J. E., Jr., 1987, *Financial Decision Making*, Rowman & Littlefield: Savage, Maryland.
- Jobson, J.D., and B.M. Korkie, 1980, Estimation for Markowitz Efficient Portfolios, *Journal of the American Statistical Association* **75**, 544-554.
- Lucas, Robert E., Jr., 1978, Asset Prices in an Exchange Economy, *Econometrica* **46**, 1429-1445.
- MacKinlay, A. Craig, 1987, On multivariate Tests of the CAPM, *Journal of Financial Economics* **18**, 341-371.
- Schwert, G. William, and Paul J. Seguin, 1990, Heteroskedasticity in Stock Returns, *The Journal of Finance* **45**, 1129-1155.
- Shanken, Jay, 1982, An Analysis of the Traditional Risk-Return Model, Ph.D. Dissertation, Carnegie-Mellon University, Pittsburgh, PA.
- Shanken, Jay, 1987, Multivariate proxies and asset pricing relations: Living with the Roll critique, *The Journal of Financial Economics* **18**, 91-110.
- Shiller, Robert J., 1982, Consumption, asset markets and macroeconomic fluctuations, *Carnegie-Rochester Series on Public Policy* **17**, 203-238.
- Shiller, Robert J., 1989, *Market volatility*, MIT Press, Cambridge, MA.
- Snow, Carl, 1991, Diagnosing Asset Pricing Models using the Distribution of Asset Returns, *Journal of Finance* **46**, 955-983.
- Tierens, Ingrid C., 1993, Pitfalls in Naive Implementations of the Hansen-Jagannathan Diagnostic, working paper, University of Chicago.

Table 1. Summary Statistics

Consumption growth is the growth rate of seasonally adjusted U.S. real per capita expenditures for consumer nondurable goods. All returns are deflated (real) returns stated as decimal fraction per period, as described in the data section. Mean is the sample mean, σ is the sample standard deviation, ρ_1 is the first order sample autocorrelation and R^2 is the R -squared value from a regression of the variable on the lagged instruments. The lagged instruments in the quarterly data set consists of a constant and the first lagged values of the consumption growth and the real return series. In the annual data set, a constant and a single lag of the two returns are used as instruments.

Variable	Mean	σ	ρ_1	R^2
Panel A: Annual Data Set: 1891-1985 (95 observations)				
Consumption Growth	0.01815	0.03470	-0.1442	0.0195
S&P 500 Real return	0.07835	0.18986	0.03646	0.0272
T-bill Real return	0.02335	0.09491	0.31871	0.1217
Panel B: Quarterly Data Set: 1964Q4-1986Q4 (93 observations)				
Consumption Growth	0.00364	0.00997	0.06570	0.0381
3-month Bill	-0.00653	0.01293	0.37726	0.21355
6-month Bill	-0.00649	0.01492	0.28553	0.19919
9-month Bill	-0.00619	0.01716	0.19880	0.17938
12-month Bill	-0.02429	0.03974	0.08040	0.07763

Table 2. Monthly Returns and Instruments

Monthly returns on 25 portfolios of common stocks are from Harvey and Kirby (1996). The portfolios are value-weighted within each industry group. The industries and their SIC codes are in the following table. *Mean* is the sample mean of the gross (one plus rate of) return, σ is the sample standard deviation and ρ_1 is the first order autocorrelation of the monthly return. R^2 is the coefficient of determination from the regression of the return on the two lagged instruments. The sample period is February of 1963 through December of 1994 (383 observations).

Industry	SIC codes	Mean	σ	ρ_1	R^2
1 Aerospace	372, 376	1.0107	0.0644	0.13	0.09414
2 Transportation	40, 45	1.0094	0.0648	0.08	0.06622
3 Banking	60	1.0086	0.0631	0.10	0.03665
4 Building materials	24, 32	1.0097	0.0608	0.09	0.06724
5 Chemicals/Plastics	281, 282, 286-289, 308	1.0094	0.0525	-0.01	0.04625
6 Construction	15-17	1.0109	0.0760	0.16	0.08692
7 Entertainment	365, 483, 484, 78	1.0135	0.0662	0.14	0.05069
8 Food/Beverages	20	1.0117	0.0449	0.05	0.03799
9 Healthcare	283, 384, 385, 80	1.0113	0.0524	0.01	0.02134
10 Industrial Mach.	351-356	1.0089	0.0587	0.05	0.06382
11 Insurance/Real Estate	63-65	1.0095	0.0581	0.15	0.05912
12 Investments	62, 67	1.0097	0.0453	0.05	0.07559
13 Metals	33	1.0075	0.0610	-0.02	0.02885
14 Mining	10, 12, 14	1.0108	0.0535	0.01	0.05654
15 Motor Vehicles	371, 551, 552	1.0095	0.0584	0.11	0.06550
16 Paper	26	1.0095	0.0536	-0.02	0.03265
17 Petroleum	13, 29	1.0102	0.0518	-0.02	0.03931
18 Printing/Publishing	27	1.0114	0.0586	0.21	0.10077
19 Professional Services	73, 87	1.0111	0.0693	0.13	0.07523
20 Retailing	53, 56, 57, 59	1.0106	0.0597	0.15	0.04893
21 Semiconductors	357, 367	1.0080	0.0559	0.08	0.07575
22 Telecommunications	366, 381, 481, 482, 489	1.0085	0.0412	-0.05	0.03498
23 Textiles/Apparel	22, 23	1.0100	0.0613	0.21	0.08511
24 Utilities	49	1.0078	0.0392	0.02	0.05663
25 Wholesaling	50, 51	1.0109	0.0614	0.13	0.03930
Dividend yield	na	1.631	0.6909	0.98	na
Treasury bill yield	na	3.8005	0.9083	0.98	na

Table 3. Finite Sample Properties of Bounds on Stochastic Discount Factor Variances

For each bound, this table shows the mean and standard deviation of the lower bound on σ_m^2 taken across 5,000 simulation trials. The bound is evaluated at $E(m)=1$. n is the effective number of assets and T is the number of times series observations. *Mean* and *Std.* refer to the unadjusted bounds, while *Adj. Mean* and *Adj. Std.* refer to the bounds adjusted for finite sample bias. The true bound is based on large-scale simulation with 1,000,000 observations.

Type of Bound		n	T	True	Mean	Std.	Adj. Mean	Adj. Std.
Panel A: Annual Data:								
fixed	bound	2	95	0.197	0.223	0.103	0.194	0.099
mult	bound	6	95	0.211	0.273	0.121	0.190	0.112
UE	bound	2	95	0.203	0.248	0.108	0.216	0.104
optimal	bound	2	95	0.212	0.265	0.115	0.232	0.110
Panel B: Quarterly Data:								
fixed	bound	4	93	0.488	0.561	0.193	0.488	0.182
mult	bound	20	93	0.914	1.564	0.414	0.992	0.319
UE	bound	4	93	0.915	1.167	0.306	1.051	0.287
optimal	bound	4	93	1.144	1.509	0.336	1.369	0.315
Panel C: Monthly Data:								
fixed	bound	25	383	0.104	0.200	0.058	0.121	0.054
mult	bound	75	383	0.313	0.626	0.114	0.305	0.092
UE	bound	25	383	0.329	0.523	0.096	0.421	0.089
optimal	bound	25	383	0.386	0.615	0.115	0.506	0.107

Table 4. Sensitivity Analyses

This table shows the true bounds, obtained by a large scale simulation with 1,000,000 observations, together with the mean and standard deviation of the bounds on σ_m^2 taken across 5,000 simulation trials. Each bound is evaluated at $E(m)=1$. *Mean* and *Std.* refer to the unadjusted bounds, while *Adj. Mean* and *Adj. Std.* refer to the bounds adjusted for finite sample bias.

Type of Bound		True	Mean	Std.	Adj. Mean	Adj. Std.
Experiment 1: Serially Dependent, Nonnormal Data						
Panel A: Annual Serially Dependent, Nonnormal Data						
fixed	bound	0.197	0.227	0.108	0.199	0.105
mult	bound	0.213	0.279	0.128	0.195	0.119
UE	bound	0.205	0.249	0.112	0.218	0.108
optimal	bound	0.211	0.271	0.123	0.239	0.118
Panel B: Quarterly Serially Dependent, Nonnormal Data						
fixed	bound	0.488	0.558	0.166	0.485	0.157
mult	bound	0.908	1.598	0.407	1.018	0.314
UE	bound	0.901	1.108	0.276	0.995	0.259
optimal	bound	1.137	1.543	0.340	1.400	0.318
Panel C: Monthly Serially Dependent, Nonnormal Data						
fixed	bound	0.104	0.203	0.061	0.123	0.057
mult	bound	0.300	0.628	0.120	0.307	0.096
UE	bound	0.321	0.509	0.098	0.408	0.091
optimal	bound	0.384	0.612	0.124	0.509	0.116

Table 4. Sensitivity Analyses (Continued)

Experiment 2: Serially Dependent, Nonnormal, Conditionally Heteroskedastic Data

Panel D: Annual Serially Dependent, Nonnormal, Conditionally Heteroskedastic Data

fixed	bound	0.201	0.233	0.112	0.204	0.108
mult	bound	0.223	0.295	0.137	0.210	0.126
UE	bound	0.213	0.264	0.117	0.232	0.112
optimal	bound	0.219	0.276	0.124	0.243	0.119

Panel E: Quarterly Serially Dependent, Nonnormal, Conditionally Heteroskedastic Data

fixed	bound	0.494	0.572	0.172	0.498	0.163
mult	bound	0.995	1.538	0.386	0.972	0.298
UE	bound	0.898	1.017	0.260	0.910	0.244
optimal	bound	1.138	1.487	0.331	1.348	0.309

Panel F: Monthly Serially Dependent, Nonnormal, Conditionally Heteroskedastic Data

fixed	bound	0.103	0.202	0.057	0.123	0.053
mult	bound	0.313	0.647	0.125	0.323	0.100
UE	bound	0.320	0.521	0.100	0.419	0.093
optimal	bound	0.381	0.584	0.104	0.478	0.096

Table 5. Evaluation of Asymptotic Standard Errors

This table evaluates the asymptotic standard errors for the fixed variance bound, adapted from Hansen, Heaton and Luttmer (1995), as in Equation (30). The variance bounds are evaluated at $E(m) = 1$. *Empirical SD* is the standard deviation of the estimated bound, taken over the 5,000 simulation trials. *Adjusted Empirical* is the standard deviation of the bias-adjusted bound. *Average Asymptotic* is the average of the asymptotic standard error across the 5,000 trials. *Adjusted Asymptotic* is the average asymptotic standard error for the bias-adjusted variance bound, adjusted for finite-sample bias according to Equation (24).

Type of artificial data	Empirical SD	Adjusted Empirical	Average Asymptotic	Adjusted Asymptotic
Panel A: Annual Data				
normal	0.103	0.099	0.100	0.097
nonnormal, homoskedastic	0.108	0.105	0.102	0.099
nonnormal, heteroskedastic	0.112	0.108	0.105	0.101
Panel B: Quarterly Data				
normal	0.193	0.182	0.174	0.165
nonnormal, homoskedastic	0.166	0.157	0.154	0.145
nonnormal, heteroskedastic	0.172	0.163	0.157	0.149
Panel C: Monthly Data				
normal	0.058	0.054	0.048	0.044
nonnormal, homoskedastic	0.061	0.057	0.050	0.047
nonnormal, heteroskedastic	0.057	0.053	0.032	0.030