

Stochastic Discount Factor Bounds with Conditioning Information

Wayne E. Ferson

Boston College

Andrew F. Siegel

University of Washington

Hansen and Jagannathan (1991) (hereafter HJ) derive restrictions on the volatility of stochastic discount factors that price a given set of returns. This article studies the sampling properties of HJ bounds that use conditioning information. One approach is to multiply the returns by the lagged variables. We also study optimized HJ bounds with conditioning information from Gallant, Hansen, and Tauchen (1990) and based on portfolios derived in Ferson and Siegel (2001). We document striking finite-sample biases in the HJ bounds, where the bounds reject asset-pricing models too often. We provide a useful bias correction. We also evaluate asymptotic standard errors for the bounds from Hansen, Heaton, and Luttmer (1995).

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 the symbol e is a vector of ones. This equation “prices” a vector of returns, R_t , on traded assets, measured as a unity plus the rates of return. The pricing is conditional on Z_{t-1} , a vector of instruments in the public information set at time $t - 1$. The random variable, m_t , that prices the assets is the *stochastic discount factor*. Different asset pricing models may be treated as different specifications for the stochastic discount factor.¹ 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.

Wayne E. Ferson is the Collins Chair in Finance, Carroll School of Management, Boston College and a Research Associate, National Bureau of Economic Research. Andrew F. Siegel is the Grant I. Butterbaugh Professor of Finance and Management Science and Adjunct Professor of Statistics, University of Washington. We are grateful to Geert Bekaert, Ravi Jagannathan, Jay Shanken, Guofu Zhou, and an anonymous referee for helpful comments. This article has also benefited from seminars at the 1998 European Finance Association Meetings, the 1998 National Bureau of Economic Research Asset Pricing Program, the 1999 American Finance Association meeting, and at Arizona State University, London Business School, Louisiana State University, University of Miami, Notre Dame, University of South Carolina, Washington University, and the University of Washington. The article was previously titled “Optimal Moment Restrictions on Stochastic Discount Factors.”

¹ See, for example, Ferson (1995) or Cochrane (2001) for reviews of how various asset pricing models specify stochastic discount factors, m_t .

Hansen and Jagannathan (1991) (hereafter HJ) 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 cannot satisfy Equation (1).

Burnside (1994) and Cecchetti, Lam, and Mark (1994) describe classic hypothesis tests based on the distance between a stochastic discount factor (SDF) and the HJ bounds. Burnside 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 and Zin (1991) model. These 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 article 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 various approaches. Given the evidence of biases in models with no conditioning information, an analysis of the sampling properties of bounds with conditioning information is useful. This article provides that analysis.

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 (2001). Our results show that the use of conditioning information in the bounds is important, and the way in which information is used is also important. When sampling error is accounted for, bounds that use no conditioning information hardly restrict the variances of stochastic discount factors to be positive, and thus have little economic content. In contrast, bounds that use conditioning information efficiently can rule out interesting SDF models.

Our article makes the following contributions. We document a strikingly large upward bias in HJ bounds with conditioning information. The biases are shown to be economically important in magnitude. For the multiplicative bounds suggested by HJ, our adjustment works very well in removing the expected finite-sample bias. The bias and the usefulness of the bias corrections are robust to reasonable changes in the data-generating process. Finally, we provide simulation evidence on the accuracy of the asymptotic standard errors for the HJ bounds, as developed in Hansen, Heaton, and Luttmer (1995). These standard errors perform well in finite samples, only mildly understating the empirical standard deviations.

The article is organized as follows. In Section 1 we review the HJ bounds with conditioning information. We present the adjustments for finite sample bias in Section 2. Section 3 describes the data. Section 4 presents some

empirical examples that illustrate the importance of the biases and our adjustments for the biases in HJ bounds. Section 5 describes our simulation study into the properties of the various methods for computing the HJ bounds. Section 6 considers the effects of alternative generating processes for the data. Section 7 evaluates the asymptotic standard errors of Hansen, Heaton, and Luttmer. Section 8 offers a summary and concluding remarks.

1. The Hansen–Jagannathan Bounds

We first consider the special case where the conditioning information is a constant, so the expectations in Equation (1) are unconditional and a stochastic discount factor is defined as any random variable m such that $E(mR) = e$. Assume that the random column n -vector R of the assets' gross returns has mean $E(R) = \mu$ and covariance matrix Σ .

Proposition 1 [HJ]. *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 HJ.

The stochastic discount factor in Equation (2) is a linear function of the asset returns, where the weights are fixed over time. We therefore refer to this case as the “fixed” bounds. HJ show that the bound is related to the maximum Sharpe ratio that can be obtained by a (fixed) 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 $[E(m)]^2$.

1.1 Bounds with conditioning information

Recent articles refine and extend the HJ bounds in several directions.² This article focuses on the use of given lagged variables, Z_{t-1} , to refine the bounds.

² HJ 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 (2000) bound the Sharpe ratios of assets' pricing errors. Hansen and Jagannathan (1997) develop measures of distance between candidate SDFs and the m that would price the assets.

Note that we take the instruments as given; thus we do not study how to choose Z_{t-1} . To understand how conditioning information can refine the HJ bounds, note that Equation (1) says

$$E(m_t R_t - e \mid Z_{t-1}) = 0. \quad (4)$$

Using conditional independence, Equation (4) is equivalent to

$$E\{(m_t R_t - e)f(Z_{t-1})\} = 0 \text{ for all bounded measurable scalar functions } f(\bullet), \quad (5)$$

where the unconditional expectation is assumed to exist. In other words, if we consider $R_t f(Z_{t-1})$ to represent the payoffs of *dynamic trading strategies*, whose prices are given by $e \cdot f(Z_{t-1})$, then the presence of the conditioning information is essentially equivalent to the condition that the SDF should price not only the original assets' payoff, but also the dynamic strategies. The larger 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, $f(\bullet)$ is a constant and the SDF must price only the original assets. HJ (1991) choose the set of functions $f(\bullet)$ to be elements of the linear function $I \otimes Z_{t-1}$, where I is the $n \times n$ identity matrix. This "multiplicative" approach has become a standard in the asset pricing literature.

We consider two additional approaches to conditioning information. The first is the *efficient portfolio* bounds, in which the set of functions $f(\bullet)$ is chosen to be the set of all portfolio weight functions. Equation (5) becomes $E\{(m_t R_t - e)'w(Z_{t-1})\} = 0$ for all n -vector valued functions $w(Z)$, with $w'(Z)e = 1$. Ferson and Siegel (2001) present closed-form solutions for the *unconditionally efficient* portfolio weights $w(Z)$, which maximize the Sharpe ratio and thus the lower bound on SDF variances.

The third approach to conditioning information is from Gallant, Hansen, and Tauchen (1990; hereafter GHT). They derive a greatest lower bound which implies that Equation (5) holds for all bounded integrable functions $f(\bullet)$. We present a closed-form expression that facilitates the implementation of their bound.³

1.2 Efficient portfolio bounds

Ferson and Siegel (2001) derive portfolios that use the given conditioning information, Z , to achieve unconditional mean-variance efficiency. These

³ 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.

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:

$$\begin{aligned}\mu_p &= E[x'(Z)\mu(Z)] \\ \sigma_p^2 &= E\{x'(Z)[\mu(Z)\mu'(Z) + \Sigma_e(Z)]x(Z)\} - \mu_p^2,\end{aligned}\tag{6}$$

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

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

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

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

where

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

Proposition 2 [Ferson and Siegel (2001)]. *Given n risky assets, the portfolio having minimum unconditional variance among portfolios with unconditional expected return μ_p is determined by the optimal weights:*

$$x'(Z) = \frac{e'\Lambda}{e'\Lambda e} + \frac{\mu_p - \alpha_2}{\alpha_3}\mu'(Z)\left(\Lambda - \frac{\Lambda ee'\Lambda}{e'\Lambda e}\right).\tag{11}$$

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

$$\sigma_p^2 = \left(\alpha_1 + \frac{\alpha_2^2}{\alpha_3}\right) - \frac{2\alpha_2}{\alpha_3}\mu_p + \frac{1 - \alpha_3}{\alpha_3}\mu_p^2.\tag{12}$$

The proof is given by Ferson and Siegel (2001).

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 Equation (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 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 Equation (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.

1.3 Optimal bounds

This section provides a convenient, closed-form expression for the optimal HJ bounds that were originally derived by GHT. 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)]:

$$\alpha(Z) = e' \Sigma_e^{-1}(Z) e$$

$$\beta(Z) = e' \Sigma_e^{-1}(Z) \mu(Z) \quad (15)$$

$$\gamma(Z) = \mu'(Z) \Sigma_e^{-1}(Z) \mu(Z)$$

Proposition 3. Optimal HJ bounds [GHT]. *The stochastic discount factor m with minimum variance for its expectation $E(m)$ that satisfies $E(mR | Z) = e$ is given by*

$$m = \zeta(Z) + [e - \zeta(Z)\mu(Z)]' \Sigma_e^{-1}(Z) [R - \mu(Z)], \quad (16)$$

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

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

A proof of Proposition 3 is available by request from the authors.⁴

⁴ 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

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.

1.4 Discussion

The optimal bounds provide the greatest lower bound on SDFs. 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 HJ bounds and the 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 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.⁵

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 and Liu (1999), is that the portfolio-based bounds remain valid when the conditional moments are incorrectly specified. Second, Ferson and Siegel (2001) 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

the m given by Equation (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 .

⁵ 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 $[E(m)]^{-1} - 1$.

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.

2. 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 return 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 (MLE) 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 Equation (19) has a noncentral chi-squared distribution, directly related to the distribution of the maximum squared Sharpe ratio, studied by Jobson and Korkie (1980). Using this distribution [also derived as a special case of Shanken (1982, 1987)] 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

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

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

Equation (21) 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 toward 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 directly applied to the multiplicative bounds of HJ. To see this, note that Equation (1) implies

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

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, $n(L+1)$. The adjustment 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 and find that the correction performs well in realistic sample sizes.

2.1 Bias correction for optimized bounds

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 estimated $\hat{\sigma}_{m^*}^2$ in the UE or optimal bounds, then an approximate bias-adjusted estimator for the bounds is*

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

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

$$m^* = E(m|Z) + [e - E(m|Z)\mu(Z)]'\Sigma_\varepsilon^{-1}(Z)[R - \mu(Z)]. \quad (25)$$

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

$$\begin{aligned} \sigma^2(m^*) &= \text{var}[E(m|Z)] \\ &+ E\{[e - E(m|Z)\mu(Z)]'\Sigma_\varepsilon^{-1}(Z)[e - E(m|Z)\mu(Z)]\}. \end{aligned} \quad (26)$$

For an estimated bound we replace $\mu(Z)$ and $\Sigma_\varepsilon^{-1}(Z)$ with their MLEs, which results in $\hat{\sigma}_{m^*}^2$. Evaluate the right-hand term of Equation (26) using iterated expectations. First, consider the conditional expectation given Z of the second term, assuming conditional joint normality of the returns, given Z , taking

$E(m | Z)$ as given and using Proposition 4. Then, taking the unconditional expectation of this result we arrive at the approximation

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

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

2.2 Implementing the bias corrections

We can rearrange Equation (24) to decompose the effects of the finite-sample bias adjustment:

$$\hat{\sigma}_{m^*, \text{adjusted}}^2 = \hat{\sigma}_{m^*}^2 - \left\{ \left(\frac{n+2}{T} \right) \hat{\sigma}_{m^*}^2 + \frac{n}{T} [E(m)]^2 \right\} + \frac{2}{T} \text{var}[E(m | Z)]. \quad (28)$$

The second term on the right side of Equation (28) may be considered a “degrees-of-freedom” adjustment, as the effect for a given $\hat{\sigma}_{m^*}^2$ and $E(m)$ depends only on n and T . This term always lowers the estimated bound, more so when n is large relative to T . The third term depends on the unconditional variance of $E(m | Z)$, and it works in the opposite direction. This term is zero in the fixed and multiplicative bounds, since $E(m | Z)$ is a constant in those bounds. We find that $(2/T)\text{var}[E(m | Z)]$ makes a very small contribution to the adjustment (less than 0.2% of the total effect) for both the UE and optimal bounds.

For the optimal bounds, $E(m | Z)$ is specified in Equation (17), and for the UE bounds

$$\text{var}[E(m_{UE} | Z)] = \left(\frac{1 - E(m)\mu_p}{\sigma_p} \right)^2 \frac{\text{var}[E(R_{UE} | Z)]}{\sigma_p^2}. \quad (29)$$

To use Equation (29) in the bias correction, we find R_{UE} , the portfolio with weights given by Equation (11), where μ_p and σ_p^2 are chosen to correspond to the point on the mean standard deviation frontier, tangent to a line drawn from $[E(m)]^{-1}$ on the expected gross return axis.⁶ Using Equations (17) and (29), consistent estimates of $\text{var}[E(m | Z)]$ are obtained from the sample variances of the fitted values of the terms that are functions of Z .

When $(2/T)\text{var}[E(m | Z)]$ is small, Equation (28) shows that the bias adjustment is essentially the degrees-of-freedom component. The adjustment

⁶ 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).

differs across the various bounds mainly when the degrees-of-freedom differ. For example, the adjustment is much larger in the multiplicative bounds, where $n(L+1)$ replaces n as the number of assets. The bias correction also depends proportionally on the magnitude of $\hat{\sigma}_m^2$ and will be smaller when the level of the bounds is lower, such as in the fixed bounds.

3. Data

We use three different datasets in our empirical illustrations. These are designed to span the range of data frequencies, sample sizes, and predictability environments most relevant to the empirical asset pricing literature. An annual and a quarterly dataset are constructed, similar to HJ. The quarterly data feature relatively strong return predictability based on interest rates, in a small cross section of bond returns. A monthly dataset provides an example representative of asset-pricing studies using a cross section of equity portfolios, with short rates and dividend yields as predictors. Summary statistics are provided in Tables 1 and 2.

The annual dataset used by HJ 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.⁷

The quarterly returns are the real, three-month holding period returns on Treasury bills with initial maturities of 3, 6, 9, and 12 months. The returns are from the Center for Research in Security Prices files for original-issue 12-month bills. Real returns are the nominal returns deflated by the component of the consumer price index 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 dataset includes the total returns (price change plus dividends) on 25 industry portfolios from Harvey and Kirby (1996), measured for the period February 1963 to December 1994.⁸ 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 & Poors 500 index.

⁷ These data are published in Shiller (1989, Tables 26.1–26.2).

⁸ We are grateful to Campbell Harvey for providing the data.

Table 1
Summary statistics

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

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 a regression of the variable on the lagged instruments. The lagged instruments in the quarterly dataset consist of a constant and the first lagged values of the consumption growth and the real return series. In the annual dataset, a constant and a single lag of the two returns are used as instruments.

Table 2
Monthly returns and instruments

	Industry	SIC codes	Mean	σ	ρ_1	R^2
1	Aerospace	372, 376	0.0107	0.0644	0.13	0.09414
2	Transportation	40, 45	0.0094	0.0648	0.08	0.06622
3	Banking	60	0.0086	0.0631	0.10	0.03665
4	Building materials	24, 32	0.0097	0.0608	0.09	0.06724
5	Chemicals/plastics	281, 282, 286–289, 308	0.0094	0.0525	−0.01	0.04625
6	Construction	15–17	0.0109	0.0760	0.16	0.08692
7	Entertainment	365, 483, 484, 78	0.0135	0.0662	0.14	0.05069
8	Food/beverages	20	0.0117	0.0449	0.05	0.03799
9	Healthcare	283, 384, 385, 80	0.0113	0.0524	0.01	0.02134
10	Industrial mach.	351–356	0.0089	0.0587	0.05	0.06382
11	Insurance/real estate	63–65	0.0095	0.0581	0.15	0.05912
12	Investments	62, 67	0.0097	0.0453	0.05	0.07559
13	Metals	33	0.0075	0.0610	−0.02	0.02885
14	Mining	10, 12, 14	0.0108	0.0535	0.01	0.05654
15	Motor vehicles	371, 551, 552	0.0095	0.0584	0.11	0.06550
16	Paper	26	0.0095	0.0536	−0.02	0.03265
17	Petroleum	13, 29	0.0102	0.0518	−0.02	0.03931
18	Printing/publishing	27	0.0114	0.0586	0.21	0.10077
19	Professional services	73, 87	0.0111	0.0693	0.13	0.07523
20	Retailing	53, 56, 57, 59	0.0106	0.0597	0.15	0.04893
21	Semiconductors	357, 367	0.0080	0.0559	0.08	0.07575
22	Telecommunications	366, 381, 481, 482, 489	0.0085	0.0412	−0.05	0.03498
23	Textiles/apparel	22, 23	0.0100	0.0613	0.21	0.08511
24	Utilities	49	0.0078	0.0392	0.02	0.05663
25	Wholesaling	50, 51	0.0109	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

Monthly returns on 25 portfolios of common stocks are from Harvey and Kirby (1996). The portfolios are value weighted within each industry group. *Mean* is the sample mean of the return, in monthly decimal fraction units, 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, which are the dividend yield and Treasury bill yield shown in the last two rows. The sample period is February 1963 through December 1994 (383 observations).

4. Empirical Examples

In this section we illustrate sample versions of the various HJ bounds. Here is how the bounds are constructed. The “fixed” bounds use no conditioning information and are determined by a fixed-weight combination of the basic asset returns, as in Equation (2). We use the maximum likelihood estimates of the mean vector and covariance matrix. When normality is not assumed, these are consistent method of moment estimates. To form the efficient portfolio and the optimal bounds we must specify the conditional mean function $\mu(Z)$ and the conditional variance function $\Sigma_e(Z)$. Here we simply regress the returns on the lagged instruments. The fitted values of the regression are our estimate of $\mu(Z)$ and the sample covariance matrix of the residuals is our estimate of the (fixed) conditional covariance matrix. The portfolio constants α_1 , α_2 , and α_3 in Equations (7)–(9) and the unconditional expectations in Equation (18) are estimated using sample averages. For the UE bounds, we use 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).

Figure 1 illustrates the various bounds with monthly datasets over the July 1963 to December 1994 period, using no finite sample adjustments. We show the volatility $\hat{\sigma}_m$ in the figures, but report the variance in subsequent tables. We summarize most of the subsequent results by reporting the values when $E(m) = 1$. The general shape of the curves, for other values of $E(m)$, are similar to those depicted here for monthly data, but the distances between the various bounds and the amount of convexity in the curves differ, depending on the amount of predictability and on the expected return-to-volatility ratios available in the datasets. For example, in the annual dataset the amount of predictability is relatively small, so the curves are closer together. In the quarterly dataset, the curves appear more convex.

A valid stochastic discount factor must lie above the bounds, “in the cup.” The bounds using conditioning information plot above the fixed bounds in Figure 1, illustrating that conditioning information allows one to rule out more stochastic discount factors. Also, there are substantial differences between the various bounds, so it matters which bound one uses.

To illustrate, we consider the example of the “three-factor model” of Fama and French (1993, 1996).⁹ Fama and French advocate a model in which three return factors describe SDF. The factors are a market portfolio return, the difference between the returns of a small-stock and a large-stock portfolio, and the difference between a high and a low book-to-market portfolio. While there is some controversy over the justification for this model, it has been popular in recent studies. If we hypothesize a multibeta pricing model, where

⁹ Using an example from Lucas’s (1978) consumption-based model, we find that the consumption SDF is so far outside the bounds (using annual, quarterly, or monthly data) that our results would not change any inferences. A simple model of habit persistence, like the one used by HJ, produces a similar result.




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

expected returns are linear in their covariances with the Fama–French factors, it implies and 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. (2002) estimate SDF formulations of the Fama–French model using a monthly dataset for July 1963 to 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 our 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 graphs. When $a(z)$ and $B(z)$ are constants and no lagged instruments are used in forming the SDF, we have an *unconditional* model (“FFU”).

Figure 1 shows that 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 for the bounds are shown, if we use the sampling variation from the simulations below, both versions of the three-factor model would be rejected using the biased bounds with conditioning information.

Note that in Figure 1 the sample multiplicative bound plots *above* the optimal bound. 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

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

Figure 2 applies our finite sample bias adjustments to the various bounds. Now the ordering of the various bounds appears reasonable, with the optimal bounds plotting above and the fixed bounds below. The effect of the adjustment to the multiplicative bounds, in particular, appears substantial.¹⁰

The conditional three-factor model SDF now plots above the fixed and multiplicative bounds, close to the UE bound, and below the optimal bounds with conditioning information. The point FFC lies 0.64 standard errors below the optimal bound, while FFU is 2.49 standard errors below.¹¹ After bias adjustment, one would not reject the conditional Fama–French SDF “FFC” using the bounds with conditioning information. This reverses the conclusion from Figure 1, illustrating that the biases in the bounds are economically significant. Furthermore, after bias adjustment one would reject FFU with the optimal bounds, but not with the multiplicative bounds. Thus the choice between bounds with conditioning information is also a matter of economic significance.

¹⁰ For some values of $E(m)$ the adjusted $\hat{\sigma}_m$ 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, and the negative adjusted bounds suggest as much. Negative values are set to zero in the figure. The local concavity of the curve for values of $\hat{\sigma}_m$ near zero is explained by the fact that the square root functions have an infinite slope at zero.

¹¹ These examples ignore the sampling error in the location of the FFC and FFU points, and should therefore be taken as illustrative. See Burnside (1994), Cecchetti, Lam, and Mark (1994), and Tierens (1993) for examples that account for this source of sampling error in specific SDF models.

5. 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 small 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 their conditional means, the simulated asset returns inherit mild serial dependence.

We estimate a first-order vector autoregression (VAR) for the lagged instruments and we use the estimated coefficient matrix as parameters of the model.¹² 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. The artificial returns are formed as the conditional means functions, 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

$$\hat{\sigma}^2(m) = a + b[E(m) - 1] + c[E(m) - 1]^2. \quad (30)$$

Thus, for each simulation trial, we record the values of a , b , and c . We use 5,000 Monte Carlo trials for each dataset. For a given $E(m)$ the values of a , b , and c , which 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 dataset. For example, in the quarterly dataset, each of the 5,000 trials uses an artificial sample of 93 quarters, representing the four assets and instruments. We conduct

¹² Such an approach may underestimate the degree of persistence in the lagged instruments [see, e.g., Bekaert, Hodrick, and Marshall (1997)]. However, we find that the artificial data generated this way match the actual sample first-order autocorrelation of the instruments to within the variability of the data across simulation trials.

simulations corresponding to the annual, quarterly, and monthly datasets. The tables summarize the results of the simulations at $E(m) = 1$.

5.1 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. (Averaging the values of $\{a, b, c\}$ across 100 simulations with 10,000 observations produces similar results.)

The true variance bounds are shown in Table 3, where the artificial data are normally distributed. The bounds are the highest in the quarterly dataset, which reflects the high Sharpe ratios that appear in samples of Treasury bill returns, consistent with HJ. Comparing the different bounds for the same dataset 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, and the optimal bound is almost four times the fixed bound. Thus, abstracting from sampling variability, both the decision to use conditioning information and the way in which the given conditioning information is used make a difference. The differences are also large in the quarterly dataset, while in the annual data, the differences between the various bounds are relatively small. In the annual dataset there is a relatively small degree of predictability, as suggested by the regression R^2 values in Table 1, so the fixed bounds are closer to the bounds with conditioning information in that case.

Table 3
Finite sample properties of bounds on stochastic discount factor variances

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	0.192	0.099
Mult bound	6	95	0.211	0.273	0.121	0.187	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	0.482	0.182
Mult bound	20	93	0.914	1.564	0.414	0.979	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	0.304	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

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 artificial data are normally distributed, homoscedastic. 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.

5.2 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; that is, the expected 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), Cecchetti, Lam, and Mark (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 two assets, and the ratio of estimated to true bounds is 113–129%. The monthly data include 383 observations on 25 assets and the ratio of estimated to true bounds is 159–200%. The quarterly data presents an intermediate case. In each dataset 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 HJ 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 MLEs 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.¹³ Since the multiplicative bounds create additional “assets,” the maximum Sharpe ratio, and thus the bound on σ_m^2 , is 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 (2001).

5.3 Bias adjustment

Table 3 reports the bounds, adjusted for finite sample bias, averaged across the 5,000 simulation trials (Adj. Mean). The average adjusted bounds range from 90% to 131% of the true bounds versus 113% to 200% before adjustment. In the annual data, all the adjusted bounds are within 10% of the true. The adjustments generally work well for the quarterly data, the optimal bound being the exception. For example, the adjustment reduces the quarterly multiplicative bound bias from 71% to 8.5%. The monthly bias is reduced

¹³ See Frost and Savarino (1988), MacKinlay (1987), and Green and Hollifield (1992) for illustrations.

from 100% to -3% . The adjustment performs quite well on the multiplicative bounds in all of the samples, and it performs even better on the fixed bounds.¹⁴

The bias adjustments provide substantial improvements in the location of the optimal and UE bounds, but they are not as accurate here as in the multiplicative case. However, before adjustment, the biases in the optimized bounds range from 28% to 59% in the quarterly and monthly data, not as severe as the multiplicative case. After adjustment the biases are roughly halved. Thus, while the bias adjustments improve the optimized bounds, they are not highly accurate. Since these are Monte Carlo simulations assuming normality, the inaccuracy of the corrections may not be attributed to the failure of the normality assumption. Thus the degrees-of-freedom adjustment, the main component of the correction, is not large enough to account for the additional complexity of the optimized bounds.

5.4 The precision of the bounds

The value of the bounds as a diagnostic tool 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 of less than 1.0. This results in smaller standard errors, and the adjusted bounds are more precise. The fixed bounds have the smallest standard deviations, and also 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 dataset. Negative values of σ_m^2 lie within two standard deviations of the true version of each bound. In the quarterly and annual datasets the bounds place substantive restrictions on σ_m^2 . The bounds with conditioning information are much more restrictive 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 dataset a value of $\sigma_m^2 = 0.128$ is two standard deviations below the true multiplicative bound. For the efficient portfolio bound, the critical value for the SDF variance is 0.148, while for the optimal bound the variance must exceed 0.173. In

¹⁴ Since we draw normally distributed data in these simulations, the adjustment to the fixed bound would be exact if the data were serially independent. We consider nonnormal and heteroscedastic artificial data below.

the quarterly sample the UE and optimal bounds have even larger advantages over the multiplicative bound. Thus when we consider both the location and the precision of the bounds, the efficient use of the conditioning information produces markedly tighter bounds. This reinforces the impression from Figure 2 that it is important to efficiently use the conditioning information in variance bounds.

6. Effects of Nonnormality and Heteroscedasticity

The previous simulations use strong assumptions about the data-generating process, including normality and homoscedasticity. Since there is evidence in the literature inconsistent with these assumptions, we explore alternative assumptions. We conduct experiments in which we generalize the data-generating process by progressively relaxing the assumptions. 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.

6.1 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 homoscedastic 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 and is not degraded by the nonnormality in the actual data.

6.2 Heteroscedasticity

Heteroscedastic data raises some new issues. First, the expected values of objects like $\mu'(Z)\Sigma_e^{-1}(Z)$, which appear in both the optimal and UE bounds, will differ from their values under homoscedasticity when the conditional

Table 4
Sensitivity analyses

Type of bound	True	Mean	Std.	Adj. mean	Adj. std.
Experiment 1: Serially dependent, nonnormal data					
Panel A: Annual 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 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 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
Experiment 2: Serially dependent, nonnormal, conditionally heteroscedastic data: Model I					
Panel D: Annual data					
Fixed bound	0.035	0.214	0.105	0.184	0.101
Mult bound	0.037	0.307	0.133	0.218	0.122
UE bound	0.036	0.257	0.111	0.225	0.106
Optimal bound	0.037	0.252	0.113	0.221	0.108
Panel E: Quarterly data					
Fixed bound	0.449	0.565	0.174	0.486	0.162
Mult bound	0.826	1.614	0.415	1.017	0.317
UE bound	0.747	1.091	0.304	0.979	0.285
Optimal bound	0.959	1.485	1.765	1.346	1.651
Panel F: Monthly data					
Fixed bound	0.050	0.200	0.057	0.121	0.053
Mult bound	0.230	0.633	0.117	0.310	0.094
UE bound	0.251	0.513	0.097	0.412	0.090
Optimal bound	0.303	0.619	0.122	0.510	0.114
Experiment 2: Serially dependent, nonnormal, conditionally heteroscedastic data: Model II					
Panel G: Annual 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 H: Quarterly 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 I: Monthly 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	0.412	0.093
Optimal bound	0.381	0.584	0.104	0.478	0.096

The true bounds are obtained by large-scale simulation with 1,000,000 observations. Each bound on σ_m^2 is evaluated at $E(m) = 1$. *Mean* and *Std.* are taken across 5,000 simulation trials. *Adj. Mean* and *Adj. Std.* refer to the bounds adjusted for finite sample bias, also with 5,000 simulation trials.

mean and elements of the conditional covariance matrix are correlated. Since this is likely to be the case, the “true” locations of the bounds may shift. Second, the issue of correctly specifying the heteroscedasticity becomes potentially important. Under heteroscedasticity, the correct data-generating 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 the next two experiments we allow for both nonnormality and conditional heteroscedasticity in the data-generating process. There are many ways to model conditional heteroscedasticity. Given the large size of the conditional covariance matrix (25×25 in the monthly data) and the fact that we embed the estimation of the model parameters in the simulation, we are motivated to use an approach to conditional heteroscedasticity that is easy to model by regression methods. To guard against conclusions that depend on the specific heteroscedasticity model, we report results for two alternative models.

Our first heteroscedasticity model, Model I, uses a single factor to generate time-varying volatility. In the annual and quarterly samples the factor, f_t , is the return of the Standard & Poor’s stock market index. In the quarterly bond return data, the factor is a short-term interest rate, the return to rolling over one-month Treasury bills.

We begin with the unexpected returns, $u_{rt} = R_t - E(R_t | Z_{t-1})$ and $u_{ft} = f_t - E(f_t | Z_{t-1})$, estimated by the regression residuals. We assume that the unexpected returns follow a factor model: $u_{rt} = u_{ft}h + \varepsilon_{rt}$, where h is a fixed n -vector of loadings and ε_{rt} has a fixed covariance matrix, $\text{cov}(\varepsilon_r)$. We then model $\text{var}(f_t | Z_{t-1}) = \text{var}(u_{ft} | Z_{t-1})$ using regression methods. Specifically, using the residuals, the $\ln(u_{ft}^2)$ are regressed on Z_{t-1} and the coefficient vector, γ , is retained as a parameter of the simulation. The conditional variance is then formed, for a given value of Z , as $\text{var}(u_{ft} | Z) = c \cdot \exp(\gamma'Z)$. (The scale factor, c , is chosen to account for Jensen’s inequality, to match the mean of $\text{var}(u_{ft} | Z_{t-1})$ to the sample variance of u_{ft} .) This leads to a model for the conditional covariance matrix, $\Sigma_\varepsilon(Z) = (hh') \text{var}(u_{ft} | Z_{t-1}) + \text{cov}(\varepsilon_r)$. Each period we generate the artificial return vector as the sum of the conditional mean vector, given the generated value of Z , plus an independent draw from the matrix of the ε_{rt} residuals, plus an independent draw from the studentized residuals, u_{ft} , which are scaled to have conditional variance $c \cdot \exp(\gamma'Z)$, and multiplied by the loading, h .

Our second heteroscedasticity model, Model II, allows each asset return to have a unique time-varying volatility, while the correlations among the returns are assumed to be fixed. The approach is essentially the same as in Model I, except the conditional standard deviations are modeled by regressing the absolute residuals, $|u_{rt}|$, on the lagged instruments, and scaling the fitted value by $\sqrt{\pi/2}$.¹⁵ The conditional covariance matrix is then formed as the

¹⁵ This approach is advocated by Davidian and Carroll (1987), and is similar to Schwert and Seguin (1990) and Ferson and Foerster (1994).

products of the conditional standard deviations and the fixed correlations. The correlations are chosen to equal the sample correlations of the regression residuals u_{it} .¹⁶

In the simulations, our artificial econometrician uses the data essentially the same way that we use it to estimate the bounds in the empirical examples and to define the data-generating process. For the fixed and multiplicative bounds, the analyst simply computes the sample estimates for the unconditional mean and covariance matrix, and with these, constructs the bounds.

Things are slightly more complicated for the optimal and UE bounds. In these bounds the functional forms of $\mu(Z)$ and $\Sigma_\varepsilon(Z)$ are taken as known, but all the parameters describing those functions must be estimated. We include the estimation of these parameters in the simulation. This represents the potential disadvantage of optimal and UE bounds, that these parameters must be estimated to implement the bounds, and their estimation creates sampling error. The constants which are the unconditional expectations of functions of $\mu(Z)$ and $\Sigma_\varepsilon(Z)$, such as in Equations (7)–(9), (17) and (18), are also estimated by the artificial econometrician.

Assuming that the analyst knows the functional forms of $\mu(Z)$ and $\Sigma_\varepsilon(Z)$, but not the true parameter values, is consistent with the previous experiments, where the analyst knows that $\mu(Z)$ is a regression function of returns on the lagged instruments and $\Sigma_\varepsilon(Z)$ is a constant. For a discussion which emphasizes the role of misspecified conditional moments in HJ bounds, see Bekaert and Liu (1999).

Panels D–F of Table 4 show the results of the heteroscedasticity experiments under Model I and panels G–I present Model II. Compared to the previous experiments, the locations of the true bounds are slightly higher in Model II, and lower in Model I. The ratios of the estimated to the true bounds, however, are similar to the previous results in Model II. The biases vary from 15% to 107% across the cases, and the relative performance of the different bounds is similar to that observed before. In Model I the biases are more extreme in percentage terms, where the true bounds are lower, but comparable in absolute magnitude for many of the cases (exceptions are the annual data, and the fixed bounds in monthly data, where the true bounds are much lower). The overall impressions are consistent with the previous experiments. 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 for the fixed and multiplicative bounds, with a few exceptions in Model I. In Model II the adjusted location of each estimated bound is within 11% of the

¹⁶ We find that this approach overstates the conditional heteroscedasticity, in the sense that the regression of the generated unexpected returns squared, on the instruments, produces a larger R^2 than the original data. We therefore shrink the conditional covariance matrices in Model II toward the fixed unconditional covariance matrix using a convex combination of the two, where the weight is selected for each dataset to match the regression R^2 with the actual data.

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.

7. Asymptotic Standard Errors

Hansen, Heaton, and Luttmer (1995; hereafter HHL) provide asymptotic distribution theory for the minimum second moment of a stochastic discount factor. This section evaluates the accuracy of the HHL standard errors in finite samples. They show that a consistent estimator for the asymptotic variance of $\hat{\sigma}_m^2$ is obtained as the asymptotic variance of $(1/T) \sum \phi_t$, with

$$\phi_t = \{-[\hat{\alpha}'(X_t - \hat{\mu})]^2 - 2\hat{\alpha}'[E(m)X_t - P_{t-1}]\}, \quad (31)$$

where $\hat{\alpha} = S^{-1}[\frac{1}{T} \sum P_{t-1} - E(m)\hat{\mu}]$. In this formulation, $X_t = R_t \otimes Z_{t-1}$ and $P_{t-1} = Z_{t-1}$ for the multiplicative bounds. For the fixed bounds, $X_t = R_t$ and $P_{t-1} = e$, the vector of ones. The sample mean is $\hat{\mu} = \frac{1}{T} \sum X_t$. We estimate the asymptotic variance of $(1/T) \sum \phi_t$ using the spectral density estimator at frequency zero from Hansen (1982). The number of autocovariance terms included is determined by examining the sample autocorrelations of ϕ_t and including the lags where the sample autocorrelations exceed two approximate standard errors.¹⁷

Table 5 summarizes the sampling properties of the HHL asymptotic standard errors for the fixed bounds. The “Unadjusted Empirical” and “Adjusted Empirical” are the standard deviations of the variance bounds from the simulations, repeated from the previous tables for convenience. “Average Unadjusted Asymptotic” is the mean value of the HHL standard errors taken over the 5,000 trials, and “Average Adjusted Asymptotic” is the mean value of HHL standard errors for the adjusted variance bounds. The table shows that the asymptotic standard errors are reliable in the annual and quarterly data. They are mildly understated, relative to the empirical standard errors, by less than 10%, in each of the four 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 homoscedastic, and sometimes more under heteroscedasticity. Overall, however, the fit is rather impressive.

Table 6 summarizes the accuracy of the HHL standard errors for the multiplicative variance bounds. We are unable to extend the analysis of HHL to cases where the conditioning information is used optimally, so this remains a

¹⁷ This criterion results in one lag in the annual and monthly datasets and four lags in the quarterly dataset.

Table 5
Evaluation of asymptotic standard errors: fixed bounds

Type of artificial data	Unadjusted empirical	Adjusted empirical	Average unadjusted asymptotic	Average adjusted asymptotic
Panel A: Annual data				
Normal	0.103	0.099	0.102	0.099
Nonnormal, homoscedastic	0.108	0.105	0.105	0.101
Nonnormal, heteroscedastic I	0.105	0.101	0.099	0.096
Nonnormal, heteroscedastic II	0.112	0.108	0.104	0.101
Panel B: Quarterly data				
Normal	0.193	0.182	0.175	0.165
Nonnormal, homoscedastic	0.166	0.157	0.154	0.146
Nonnormal, heteroscedastic I	0.174	0.162	0.157	0.148
Nonnormal, heteroscedastic II	0.172	0.163	0.157	0.149
Panel C: Monthly data				
Normal	0.058	0.054	0.048	0.044
Nonnormal, homoscedastic	0.061	0.057	0.050	0.047
Nonnormal, heteroscedastic I	0.057	0.053	0.049	0.046
Nonnormal, heteroscedastic II	0.057	0.053	0.032	0.030

This table evaluates the asymptotic standard errors for the fixed variance bound, from Hansen, Heaton, and Luttmer (1995), as given by Equation (31). The variance bounds are evaluated at $E(m) = 1$. *Unadjusted empirical* 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 unadjusted asymptotic* is the average of the asymptotic standard error across the 5,000 trials. *Average adjusted asymptotic* is the average asymptotic standard error for the bias-adjusted variance bound, adjusted for finite-sample bias according to Equation (24).

Table 6
Evaluation of asymptotic standard errors: multiplicative bounds

Type of artificial data	Unadjusted empirical	Adjusted empirical	Average unadjusted asymptotic	Average adjusted asymptotic
Panel A: Annual data				
Normal	0.121	0.112	0.119	0.113
Nonnormal, homoscedastic	0.128	0.119	0.125	0.118
Nonnormal, heteroscedastic I	0.133	0.122	0.127	0.120
Nonnormal, heteroscedastic II	0.137	0.126	0.125	0.118
Panel B: Quarterly data				
Normal	0.414	0.319	0.370	0.285
Nonnormal, homoscedastic	0.407	0.314	0.365	0.282
Nonnormal, heteroscedastic I	0.415	0.317	0.370	0.286
Nonnormal, heteroscedastic II	0.386	0.298	0.353	0.272
Panel C: Monthly data				
Normal	0.114	0.092	0.091	0.079
Nonnormal, homoscedastic	0.120	0.096	0.093	0.081
Nonnormal, heteroscedastic I	0.117	0.114	0.094	0.081
Nonnormal, heteroscedastic II	0.125	0.100	0.038	0.028

This table evaluates the asymptotic standard errors for the multiplicative variance bound, from Hansen, Heaton, and Luttmer (1995), as given by Equation (31). The variance bounds are evaluated at $E(m) = 1$. *Unadjusted empirical* 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 unadjusted asymptotic* is the average of the asymptotic standard error across the 5,000 trials. *Average adjusted asymptotic* is the average asymptotic standard error for the bias-adjusted variance bound, adjusted for finite-sample bias according to Equation (24).

topic for future research. The results for the multiplicative bounds are similar to Table 5. In the annual and quarterly data the standard errors are mildly understated, usually by less than 10%. In the monthly data the understatement is more severe, often exceeding 20%.

8. Summary and Conclusions

This article offers a number of refinements and insights into the volatility bounds on SDFs, first developed by HJ. 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 “multiplicative” 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 and show with simulations that it works very well in controlling the bias in the multiplicative bounds.

We compare Hansen and Jagannathan’s “multiplicative” approach with two alternative approaches to the use of conditioning information. One approach, following Ferson and Siegel (2001), is based on unconditional mean variance efficient portfolio strategies in the presence of conditioning information. We call these the “efficient portfolio” (UE) bounds. The second approach, based on GHT, 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 HHL. 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 large, relative to the other bounds with conditioning information. As a result, the multiplicative bounds are less restrictive of SDF variances once their sampling error is taken into account. If we could use only one version of the bounds with conditioning information, based on these results, it would not be the multiplicative bound.

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. Accounting for sampling error, the optimal bounds

are the most restrictive of SDF variances. However, the bias adjustment is the least effective for these bounds, and they are not robust to the specification of the conditional moments. 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 data-generating process, and where the most restrictive bound is desired.

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, and smaller bias prior to any adjustments. 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 to the specification of the conditional moments and precision of the bounds are the important concerns.

4. *HHL asymptotic standard errors* for the fixed and multiplicative bounds are reasonably accurate, mildly understated in our annual and quarterly samples by less than 10%. The understatement is worse when the number of assets is large relative to the number of time-series observations. The asymptotic standard errors are easy to compute and should be useful in similar applications.

References

- Bansal, R., and B. N. Lehmann, 1997, "Growth-Optimal Restrictions on Asset Pricing Models," *Macroeconomic Dynamics*, 1, 1–22.
- Bekaert, G., and J. Liu, 1999, "Conditioning Information and Variance Bounds on Pricing Kernels," working paper, National Bureau of Economic Research.
- Bekaert, B., R. J. Hodrick, and D. A. Marshall, 1997, "On Biases in Tests of the Expectations Hypothesis of the Term Structure of Interest Rates," *Journal of Financial Economics*, 44, 309–348.
- Balduzzi, P., and H. Kallal, 1997, "Risk Premia and Variance Bounds," *Journal of Finance*, 52, 1913–1949.
- Burnside, C., 1994, "Hansen-Jagannathan Bounds as Classical Tests of Asset-Pricing Models," *Journal of Business and Economic Statistics*, 12, 57–79.
- Cecchetti, S. G., P. Lam, and N. Mark, 1994, "Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns," *Journal of Finance*, 59, 123–152.
- Cochrane, J. H., 1996, "A Cross-Sectional Test of a Production-Based Asset Pricing Model," *Journal of Political Economy*, 104, 572–621.
- Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, NJ.
- Cochrane, J. H., and L. P. Hansen, 1992, "Asset Pricing Explorations for Macroeconomists," *NBER Macroeconomic Annual*, 7, 115–165.
- Cochrane, J. H., and J. Saa'-Requejo, 2000, "Beyond Arbitrage: Good-Deal Pricing of Derivatives in Incomplete Markets," *Journal of Political Economy*, 108, 79–119.
- Davidian, M., and R. J. Carroll, 1987, "Variance Function Estimation," *Journal of the American Statistical Association*, 82, 1079–1091.

- Dybvig, P., and J. Ingersoll, 1982, "Mean Variance Theory in Complete Markets," *Journal of Business*, 55, 233–251.
- Efron, B., 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, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.
- Fama, E. F., and K. R. French, 1996, "Multifactor Explanations of Asset Pricing Anomalies," *Journal of Finance*, 51, 55–87.
- Farnsworth, H., W. Ferson, D. Jackson, and S. Todd, 2002, "Performance Evaluation with Stochastic Discount Factors," forthcoming in *Journal of Business*.
- Ferson, W. E., 1995, "Theory and Empirical Testing of Asset Pricing Models," *Handbooks in OR & MS*, Vol. 9, Elsevier, Amsterdam.
- Ferson, W. E., and S. 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, W. E., and C. R. Harvey, 1992, "Seasonality and Consumption-Based Asset Pricing Models," *Journal of Finance*, 47, 511–552.
- Ferson, W. E., and R. Jagannathan, 1996, "Econometric Evaluation of Asset Pricing Models," in G. S. Maddala and C. R. Rao (eds.), *Handbook of Statistics*, Vol. 14, Statistical Methods in Finance, Amsterdam, North Holland.
- Ferson, W. E., and A. F. Siegel, 2001, "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56, 967–982.
- Frost, P., and J. Savarino, 1988, "For Better Performance: Constrain Portfolio Weights," *Journal of Portfolio Management*, 15, 29–34.
- Gallant, A. R., L. P. Hansen, and G. Tauchen, 1990, "Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution," *Journal of Econometrics*, 45, 141–179.
- Green, R., and B. Hollifield, 1992, "When Will Mean Variance Efficient Portfolios be Well Diversified?" *Journal of Finance*, 47, 1785–1809.
- Hansen, L. P., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.
- Hansen, L. P., J. Heaton, and E. Luttmer, 1995, "Econometric Evaluation of Asset Pricing Models," *Review of Financial Studies*, 8, 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 R. Jagannathan, 1997, "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.
- 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, 587–613.
- Harvey, C. R., and C. Kirby, 1996, "Analytic Tests of Factor Pricing Models," unpublished working paper, Duke University.
- Ingersoll, J. E., Jr., 1987, *Financial Decision Making*, Rowman & Littlefield, Savage, MD.
- Jobson, J. D., and B. M. Korkie, 1980, "Estimation for Markowitz Efficient Portfolios," *Journal of the American Statistical Association*, 75, 544–554.

- Lucas, R. E., Jr., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.
- MacKinlay, A. C., 1987, "On Multivariate Tests of the CAPM," *Journal of Financial Economics*, 18, 341–371.
- Schwert, G. W., and P. J. Seguin, 1990, "Heteroskedasticity in Stock Returns," *Journal of Finance*, 45, 1129–1155.
- Shanken, J., 1982, "An Analysis of the Traditional Risk-Return Model," Ph.D. dissertation, Carnegie-Mellon University, Pittsburgh, PA.
- Shanken, J., 1987, "Multivariate Proxies and Asset Pricing Relations: Living with the Roll Critique," *Journal of Financial Economics*, 18, 91–110.
- Shiller, R. J., 1982, "Consumption, Asset Markets and Macroeconomic Fluctuations," *Carnegie-Rochester Series on Public Policy*, 17, 203–238.
- Shiller, R. J., 1989, *Market Volatility*, MIT Press, Cambridge, MA.
- Snow, C., 1991, "Diagnosing Asset Pricing Models Using the Distribution of Asset Returns," *Journal of Finance*, 46, 955–983.
- Tierens, I. C., 1993, "Pitfalls in Naive Implementations of the Hansen-Jagannathan Diagnostic," working paper, University of Chicago.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.