

Kan, Raymond; Robotti, Cesare

Working Paper

The exact distribution of the Hansen-Jagannathan bound

Working Paper, No. 2008-9

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Kan, Raymond; Robotti, Cesare (2008) : The exact distribution of the Hansen-Jagannathan bound, Working Paper, No. 2008-9, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/70721>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Exact Distribution of the
Hansen-Jagannathan Bound

Raymond Kan and Cesare Robotti

Working Paper 2008-9
February 2008

The Exact Distribution of the Hansen-Jagannathan Bound

Raymond Kan and Cesare Robotti

Working Paper 2008-9

February 2008

Abstract: Under the assumption of multivariate normality of asset returns, this paper presents a geometrical interpretation and the finite-sample distributions of the sample Hansen-Jagannathan (1991) bounds on the variance of admissible stochastic discount factors, with and without the nonnegativity constraint on the stochastic discount factors. In addition, since the sample Hansen-Jagannathan bounds can be very volatile, we propose a simple method to construct confidence intervals for the population Hansen-Jagannathan bounds. Finally, we show that the analytical results in the paper are robust to departures from the normality assumption.

JEL classification: G12

Key words: Hansen-Jagannathan bound, exact distribution, no-arbitrage

Raymond Kan gratefully acknowledges financial support from the National Bank Financial of Canada. The authors also thank Wayne Ferson, Sergei Sarkissian, Tim Simin, Chu Zhang, Guofu Zhou, and seminar participants at the Federal Reserve Bank of Atlanta for helpful comments. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Raymond Kan, Joseph L. Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, Ontario, Canada M5S 3E6, 416-978-4291, kan@chass.utoronto.ca, or Cesare Robotti (contact author), Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, 404-498-8543, cesare.robotti@atl.frb.org.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's Web site at www.frbatlanta.org. Click "Publications" and then "Working Papers." Use the WebScriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

The Exact Distribution of the Hansen-Jagannathan Bound

Under the law of one price, Hansen and Jagannathan (1991) derive a lower volatility bound (unconstrained HJ-bound hereafter) that every valid stochastic discount factor (SDF) must satisfy. In addition, Hansen and Jagannathan (1991) propose a tighter volatility bound (constrained HJ-bound hereafter) that is applicable to nonnegative SDFs. The unconstrained HJ-bound has received wide attention in the literature. Examples include Snow (1991), Bekaert and Hodrick (1992), Ferson and Harvey (1992), Backus, Gregory, and Telmer (1993), Cecchetti, Lam, and Mark (1994), Burnside (1994), Heaton (1995), and Epstein and Zin (2001), among many others. In addition, Ferson and Siegel (2003) and Bekaert and Liu (2004) show how conditioning information can be used to optimally tighten the unconstrained HJ-bound; while Kan and Zhou (2006) tighten the unconstrained HJ-bound by making the SDF explicitly a function of a set of state variables. Although the constrained HJ-bound is sharper than the unconstrained HJ-bound and is theoretically appealing, the constrained HJ-bound has not received nearly as much attention as the unconstrained HJ-bound in empirical work. The few empirical papers that use the constrained bound besides Hansen and Jagannathan (1991) are Cecchetti, Lam, and Mark (1994), Burnside (1994), He and Modest (1995), Balduzzi and Kallal (1997), and Hagiwara and Herce (1997).

In this paper, we provide a geometrical interpretation and the finite-sample distributions of both the unconstrained and constrained HJ-bounds.¹ While there is a well-known mapping between the unconstrained HJ-bound and the mean-variance frontier of portfolio returns, the mapping between the constrained HJ-bound and the mean-variance frontier that we provide in this paper is new. We show that the linkage between the unconstrained HJ-bound and the mean-variance frontier also exists for the case of the constrained HJ-bound, except that we need to replace the mean and variance of the portfolio returns by the truncated mean and truncated variance of portfolio returns.

As we mentioned above, the constrained HJ-bound has not been very popular in the literature. We suspect that the lack of popularity of the constrained HJ-bound is due to its computational difficulty. This is because when there are N assets, one has to solve N nonlinear equations in order to obtain the constrained HJ-bound. In this paper, we show that under the assumption that

¹Hansen, Heaton, and Luttmer (1995) provide the asymptotic distributions of the unconstrained and constrained sample HJ-bounds.

returns are multivariate normally distributed, the constrained HJ-bound has a very simple analytical expression. This analytical expression allows us to obtain a maximum likelihood estimator of the constrained HJ-bound which is simpler and more precise than the traditional nonparametric estimator of the constrained HJ-bound. In addition, we provide an approximate unbiased estimator of the constrained HJ-bound with improved finite sample properties.

As documented by Burnside (1994), Cecchetti, Lam, and Mark (1994), and Ferson and Siegel (2003), the sample HJ-bounds can have a large finite sample upward bias. Although Ferson and Siegel (2003) provide a bias adjustment for the sample unconstrained HJ-bound, the adjusted estimator can still be very volatile. In this paper, we present the exact distributions of the unconstrained and constrained sample HJ-bounds under the multivariate normality assumption. In addition, we show that under general distributional assumptions, the traditional nonparametric estimator of the constrained HJ-bound does not have any finite moment. Finally, we propose a simple method to construct confidence intervals for the unconstrained and constrained HJ-bounds.

An example may help to illustrate the importance of reporting confidence intervals for the HJ-bounds instead of only presenting their point estimates. In Figure 1, we provide a characterization of the equity premium puzzle using Shiller's (1989) data.

Figure 1 about here

The sample unconstrained HJ-bound of stocks and bonds is represented by the solid line. In addition, we plot the mean and standard deviation of the SDF implied by a representative consumer model with time-separable power utility function of the form

$$U(c_t) = \frac{c_t^{1-\gamma} - 1}{1-\gamma}, \quad (1)$$

for values of γ ranging between zero and 20 (in increments of one), where $\gamma \geq 0$ is the parameter of relative risk aversion and c_t is consumption at time t .² The SDF implied by this type of preferences is given by

$$m_t = \beta \left(\frac{c_t}{c_{t-1}} \right)^{-\gamma}, \quad (2)$$

²The data are annual (from 1891 to 1985) and are obtained from Tables 26.1 and 26.2 of Shiller (1989). The asset returns are the real returns on the S&P composite stock price index and commercial paper. Consumption growth is the growth rate of seasonally adjusted U.S. real per capita expenditures for consumer nondurable goods.

where β , the subjective discount factor, is set equal to 0.99.³

The point estimate of the unconstrained HJ-bound leads us to conclude that the SDF is rejected by the data for values of γ of 18 or less. Only for large values of γ , the SDF lies above the sample HJ-bound and qualifies as a valid SDF. In Figure 1, we also plot the bias-adjusted sample unconstrained HJ-bound (dashed line) of Ferson and Siegel (2003). Since the bias adjustment is very small when there are only two assets, we arrive at the same conclusion even when we use the bias-adjusted version of the sample HJ-bound. As pointed out by Gregory and Smith (1992) and Burnside (1994), such a comparison ignores the sampling error in the sample HJ-bound and can lead to a false rejection of the model. To get an idea of where the population HJ-bound may actually lie, we construct 95% confidence intervals for the unconstrained HJ-bound (dotted lines) using the methodology described later in the paper. The confidence intervals in Figure 1 are quite wide, indicating that there is substantial uncertainty about the exact location of the HJ-bound. Taking sampling error into account, the SDF in (2) might not be entirely at odds with the data even for small values of γ . In fact, the SDF lies inside the 95% confidence intervals for the unconstrained HJ-bound for values of γ of one and two.⁴ This stands in sharp contrast to the strong rejection of the model that we obtain when merely relying on the sample HJ-bound.

Since the SDF in (2) is positive, it seems reasonable to compare it with the more demanding HJ-bound that imposes the nonnegativity constraint. In Figure 2, we provide a comparison of this SDF with the constrained HJ-bound. We present three different estimators of the sample constrained HJ-bound. The dashed-dotted line represents the nonparametric estimator of the constrained HJ-bound that is used in the existing literature whereas the solid and dashed lines represent, in the order, the new maximum likelihood and approximate unbiased estimators of the constrained HJ-bound that we develop in this paper. For this particular example, all three estimators of the

³Since we consider Figure 1 to be primarily diagnostic, we do not present confidence regions for the mean-standard deviation pairs of the SDF or formally test whether the point estimates of the SDF lie outside the HJ-bound. Burnside (1994) and Cecchetti, Lam, and Mark (1994) develop classical hypothesis tests based on the distance between a given SDF and the HJ-bound. These studies find that the point estimates of the SDF plot outside the sample unconstrained and constrained HJ-bound too often when the model is true. Otrok, Ravikumar, and Whiteman (2002) use Monte Carlo simulations to derive finite-sample critical values of the test statistics developed by Burnside (1994) and Cecchetti, Lam, and Mark (1994).

⁴Instead of constructing confidence intervals for the HJ-bound, Burnside (1994) constructs confidence regions for β and γ . He finds that the 95% confidence regions for the (β, γ) pairs contain part of the parameter space described by Mehra and Prescott (1985) as “reasonable.” Burnside’s approach takes into account the sampling variability of the SDF, whereas our approach is only concerned with the variability of the sample HJ-bound. The advantage of our approach is that the confidence intervals for the HJ-bound are computed based on the exact distribution. In addition, once the confidence intervals for the HJ-bound are computed, they can be used for evaluating multiple SDFs.

constrained HJ-bounds are quite close to each other and lead us to conclude that the SDF is rejected by the data for values of γ of 18 or less. However, since the point estimates of the sample constrained HJ-bounds are very volatile, we also present the 95% confidence intervals for the constrained HJ-bound using the dotted lines. Just like in the case of the unconstrained HJ-bound, the SDF lies inside the 95% confidence intervals for the constrained HJ-bound for reasonable values of γ of one and two.

Figure 2 about here

The rest of the paper is organized as follows. The next section presents the unconstrained and constrained HJ-bounds and our main results on the constrained HJ-bound under the assumption that returns are multivariate normally distributed. In Section 2, we summarize the asymptotic distributions of the sample unconstrained and constrained HJ-bounds and present a new maximum likelihood estimator of the constrained HJ-bound. In section 3, under the multivariate normality assumption, we present the finite-sample distributions of the sample unconstrained and constrained HJ-bounds and derive an approximate unbiased estimator of the constrained HJ-bound. In addition, we present a method for constructing exact confidence intervals for the unconstrained and constrained HJ-bounds. Finally, we investigate the robustness of our finite sample results to departures from normality. The last section summarizes our findings and the Appendix contains proofs of all propositions.

1. Hansen-Jagannathan Bounds

In this section, we summarize existing results and present some new ones on the HJ-bounds. This section is divided into three subsections. In subsection 1, we present the unconstrained HJ-bound for the case in which we only require the SDFs to satisfy the law of one price. In subsection 2, we present the constrained HJ-bound for the case in which we also impose the nonnegativity constraint on the SDFs. Most of the results in these two subsections are well known from the work of Hansen and Jagannathan (1991). In subsection 3, we present the constrained HJ-bound under the assumption that returns are multivariate normally distributed. While the normality assumption is restrictive, it allows us to (1) understand the determinants of the difference between the unconstrained and

constrained HJ-bounds, (2) establish a connection between the minimum-variance frontier and the constrained HJ-bound; and (3) conduct finite sample inference on the sample HJ-bounds.

The basic setup is as follows. Denote the vector of gross returns on the N risky assets by R and the mean and the covariance matrix of R by $\mu = E[R]$ and $V = \text{Var}[R]$, respectively.⁵ In addition, we assume that the gross risk-free rate is R_0 , so that there are altogether $N + 1$ assets. In some applications, there is no risk-free asset and R_0 will be treated as a free variable. The HJ-bound will then be expressed as a function of R_0 .

The analyses for both the constrained and the unconstrained cases are very similar. We first start off with an optimal portfolio problem and then write the SDF as a function of the gross return on the optimal portfolio. The variance of this SDF gives us the HJ-bound.

1.1 Unconstrained Hansen-Jagannathan Bound

When the law of one price holds, there exists an SDF m that prices all the risky assets correctly

$$E[mR] = 1_N, \tag{3}$$

where 1_N is an N -vector of ones. In addition, the risk-free rate R_0 restricts the mean of m because $E[mR_0] = 1 \Rightarrow \mu_m \equiv E[m] = 1/R_0$. There can be many m 's that price the $N + 1$ assets correctly, but we are interested in finding the one that has the lowest variance. Instead of directly solving this problem, Hansen and Jagannathan (1991) propose to solve a dual problem. The dual problem consists in finding a portfolio that minimizes the second moment of its gross return. Denote by w the portfolio weights in the N risky assets and by $1 - w'1_N$ the portfolio weight in the risk-free asset. The gross return on the portfolio is given by

$$R_p = (1 - w'1_N)R_0 + w'R = R_0 + w'(R - R_01_N). \tag{4}$$

The portfolio that minimizes the second moment is the solution to the following problem

$$\min_w E[(R_0 + w'(R - R_01_N))^2]. \tag{5}$$

⁵Although we assume R to be gross returns, we can easily change the setup to allow for some or all of R to be excess returns (i.e., returns on zero investment portfolios). All we need to do is to replace 1_N with q in our subsequent analysis, where q is the vector of the costs of the N risky assets (with elements of zero or one to indicate whether the returns are excess returns or gross returns).

Denoting the minimum second moment portfolio by p^* , it is straightforward to show that its weights in the N risky assets are given by

$$w^* = -\frac{R_0}{1 + \theta_0^2} V^{-1}(\mu - R_0 \mathbf{1}_N), \quad (6)$$

where

$$\theta_0^2 = (\mu - R_0 \mathbf{1}_N)' V^{-1}(\mu - R_0 \mathbf{1}_N) = a - 2bR_0 + cR_0^2, \quad (7)$$

and $a = \mu' V^{-1} \mu$, $b = \mathbf{1}'_N V^{-1} \mu$, and $c = \mathbf{1}'_N V^{-1} \mathbf{1}_N$ are the three efficiency set constants that characterize the minimum-variance frontier of the N risky assets. Note that w^* is proportional to the weights of the tangency portfolio (i.e., the portfolio that maximizes the Sharpe ratio) which has weights $V^{-1}(\mu - R_0 \mathbf{1}_N)/(b - cR_0)$ in the risky assets. This suggests that p^* is a linear combination of the risk-free asset and the tangency portfolio and its gross return is given by

$$R_{p^*} = R_0 + w^{*'}(R - R_0 \mathbf{1}_N). \quad (8)$$

It is easy to verify that $\mu_{p^*} \equiv E[R_{p^*}] = R_0/(1 + \theta_0^2)$, $\sigma_{p^*}^2 \equiv \text{Var}[R_{p^*}] = R_0^2 \theta_0^2 / (1 + \theta_0^2)^2$, $E[R_{p^*}^2] = R_0^2 / (1 + \theta_0^2)$, and the squared Sharpe ratio of p^* is θ_0^2 . Define an SDF as

$$m_0 = \frac{R_{p^*}}{\|R_{p^*}\|^2} = \frac{1 - (\mu - R_0 \mathbf{1}_N)' V^{-1}(R - \mu)}{R_0}, \quad (9)$$

where $\|X\| = E[X^2]^{\frac{1}{2}}$. Lemma 1 summarizes the properties of m_0 that are given in Hansen and Jagannathan (1991).

Lemma 1. *For m_0 defined in (9), we have (1) $E[m_0] = 1/R_0$, (2) $E[m_0 R] = \mathbf{1}_N$, (3) $\|m_0\| = 1/\|R_{p^*}\|$, (4) $R_{p^*} = m_0/\|m_0\|^2$, (5) for any admissible SDF m with $E[m] = 1/R_0$, we have $\text{Cov}[m, m_0] = \text{Var}[m_0]$ and $\text{Var}[m] \geq \text{Var}[m_0]$.*

The first two properties tell us that m_0 is indeed a valid SDF that correctly prices the risk-free asset and the N risky assets. The third and the fourth properties show the duality between m_0 and R_{p^*} . The last property suggests that $\text{Var}[m_0]$ provides a lower bound for the variance of all admissible SDFs with $E[m] = 1/R_0$. It is straightforward to show that $\text{Var}[m_0]$ is given by

$$\text{Var}[m_0] = \frac{(\mu - R_0 \mathbf{1}_N)' V^{-1}(\mu - R_0 \mathbf{1}_N)}{R_0^2} = a\mu_m^2 - 2b\mu_m + c \equiv \sigma_0^2, \quad (10)$$

where $\mu_m = 1/R_0$ and σ_0^2 is called the unconstrained HJ-bound. Since every admissible SDF must be at least as volatile as m_0 , σ_0^2 can be used as a model diagnostic for a proposed SDF. Note that

σ_0^2 is a quadratic function of μ_m and it only depends on μ and V through the three efficiency set constants a , b , and c .

Hansen and Jagannathan (1991) provide a linkage between the minimum-variance frontier and the unconstrained HJ-bound. The basic relation is the following

$$\theta_0 = \frac{\sigma_{p^*}}{\mu_{p^*}} = \frac{\sigma_0}{\mu_m}. \quad (11)$$

In Figure 3, we provide a graphical illustration of this relation in the space of (σ_p, μ_p) . When the risk-free rate is R_0 , the two straight lines emanating from the point $(0, R_0)$ represent the minimum-variance frontier of the risk-free asset and the N risky assets. These two straight lines have a slope of θ_0 which is equal to the absolute value of the Sharpe ratio of the tangency portfolio. Since $E[R_p^2] = \mu_p^2 + \sigma_p^2$, the portfolio with minimum second moment has the shortest distance from the origin. In order to locate the portfolio with minimum second moment, we draw a circle with its center at the origin, and the location of the minimum second moment portfolio p^* can be obtained from the point where the circle is tangent to the minimum-variance frontier of the risk-free and risky assets. Suppose that we draw a solid line joining the origin and p^* and a horizontal dotted line at the level of $1/R_0 = \mu_m$. Since σ_{p^*}/μ_{p^*} is also equal to θ_0 , the intersection point of these two lines has a horizontal distance of $\mu_m\theta_0 = \sigma_0$ from the y -axis. Therefore, this distance gives us the HJ-bound on the standard deviation of admissible SDFs.

Figure 3 about here

1.2 Constrained Hansen-Jagannathan Bound

Hansen and Jagannathan (1991) suggest that for evaluating SDFs that are nonnegative, we can tighten the HJ-bound by imposing a nonnegativity constraint on the admissible SDFs. In order to find a nonnegative minimum-variance SDF, Hansen and Jagannathan (1991) propose to first solve the following dual portfolio problem

$$\min_w E[\max[0, R_0 + w'(R - R_0\mathbf{1}_N)]^2]. \quad (12)$$

The problem amounts to finding a portfolio that has minimum truncated second moment, i.e., $\min_w E[R_p^{+2}]$, where $R_p^+ = \max[0, R_p]$. This portfolio problem is nontrivial to solve because the truncated second moment of a portfolio depends on the joint distribution of the returns on the

risky assets. Therefore, unless the joint distribution of R is completely characterized by its first two moments (e.g., multivariate elliptical distribution), knowing the mean and the covariance matrix of R is in general not sufficient for us to solve this problem.

Denote the portfolio with minimum truncated second moment by q^* . Hansen and Jagannathan (1991) show that we can use the gross return on this portfolio to construct a nonnegative SDF

$$m_c = \frac{R_{q^*}^+}{\|R_{q^*}^+\|^2}. \quad (13)$$

Lemma 2 summarizes the properties of m_c .

Lemma 2. *For m_c defined in (13), we have (1) $E[m_c] = 1/R_0$, (2) $E[m_c R] = 1_N$, (3) $\|m_c\| = 1/\|R_{q^*}^+\|$, (4) $R_{q^*}^+ = m_c/\|m_c\|^2$, (5) for any admissible SDF m with $E[m] = 1/R_0$ and $m \geq 0$, we have $\text{Cov}[m, m_c] \geq \text{Var}[m_c]$ and $\text{Var}[m] \geq \text{Var}[m_c]$.*

The first two properties tell us that m_c is indeed a valid SDF that correctly prices the risk-free asset and the N risky assets. The third and the fourth properties show the duality between m_c and $R_{q^*}^+$. The last property suggests that $\text{Var}[m_c]$ provides a lower bound for the variance of all admissible SDFs that are nonnegative. In many ways, Lemma 2 is almost identical to Lemma 1 after we replace m_0 by m_c and R_{p^*} by $R_{q^*}^+$. The only difference is in the last property. In Lemma 1, we have $\text{Cov}[m, m_0] = \text{Var}[m_0]$ but in Lemma 2 we only have an inequality of $\text{Cov}[m, m_c] \geq \text{Var}[m_c]$. The reason is that $R_{q^*}^+$ is not a portfolio return. As a result, we do not have the exact pricing result of $E[m R_{q^*}^+] = 1$, but just the inequality of $E[m R_{q^*}^+] \geq 1$.

There is a well known mapping between the unconstrained HJ-bound and the mean-variance frontier of portfolio returns as given in (11). As it turns out, we can develop a similar mapping between the constrained HJ-bound and the truncated mean-variance frontier. Let $\mu_p^+ \equiv E[R_p^+]$ and $\sigma_p^+ \equiv (\text{Var}[R_p^+])^{\frac{1}{2}}$ be the truncated mean and standard deviation of the gross return on portfolio p , respectively. The following lemma presents the linkage between the truncated mean-variance frontier and the constrained HJ-bound.

Lemma 3. *Define the squared truncated Sharpe ratio of portfolio q^* as*

$$\theta_c^2 = \frac{(E[R_{q^*}^+] - R_0)^2}{\text{Var}[R_{q^*}^+]}. \quad (14)$$

We have

$$\theta_c = \frac{\sigma_{q^*}^+}{\mu_{q^*}^+} = \frac{\sigma_c}{\mu_m}. \quad (15)$$

Comparing (15) with (11), we see that the linkage between the unconstrained HJ-bound and the mean-variance frontier also exists for the case of the constrained HJ-bound, except that we need to replace the mean and variance by the truncated mean and truncated variance.

Although the results on the constrained and unconstrained HJ-bounds are quite similar, the constrained HJ-bound is seldom used in the empirical literature. We believe part of the reason is that it is difficult to solve the minimum truncated second moment portfolio problem in (12). Besides depending on the joint distribution of R , this problem is also highly nonlinear and there is generally no closed-form solution. Without an analytical solution, it is difficult for researchers to understand what is the portfolio q^* that minimizes the truncated second moment. As a result, we cannot plot the truncated mean-variance frontier and it is hard to visualize the relation between the truncated minimum-variance frontier and the constrained HJ-bound as described in Lemma 3. In order to overcome these problems, we need to make stronger assumptions on the joint distribution of R . In the next subsection, we make the additional assumption that the returns are multivariate normally distributed. With this assumption, we can obtain an analytical solution to the minimum truncated second moment portfolio problem.⁶ While returns are certainly not normal, we view this as a good working approximation for monthly and annual returns. More importantly, the normality assumption allows us to obtain a better understanding of the constrained HJ-bound that is hard to come by under more general distributional assumptions.

1.3 Constrained Hansen-Jagannathan Bound under Normality

We assume that $R \sim N(\mu, V)$. With the multivariate normality assumption on R , the portfolio return $R_p = R_0 + w'(R - R_0 1_N)$ is also normally distributed with mean and variance given by $\mu_p = R_0 + w'(\mu - R_0 1_N)$ and $\sigma_p^2 = w'Vw$, respectively. Lemma 4 presents the formulae for the truncated first and second moments of R_p which is the first step towards solving the minimum

⁶Cecchetti, Lam, and Mark (1992) also study the constrained HJ-bound under the normality assumption. However, their analysis does not lead to a closed-form solution of the constrained HJ-bound. In addition to the multivariate normality case, we also obtain an analytical solution for the constrained HJ-bound under the multivariate t -distribution assumption, and the results are available upon request. It is also possible to generalize our analysis to the larger class of multivariate elliptical distributions.

truncated second moment portfolio problem.

Lemma 4. *Suppose that the gross return on a portfolio, R_p , is normally distributed with mean μ_p and variance σ_p^2 . The truncated first and second moments of R_p are given by*

$$E[R_p^+] = \mu_p \Phi(\eta) + \sigma_p \phi(\eta), \quad (16)$$

$$E[R_p^{+2}] = (\mu_p^2 + \sigma_p^2) \Phi(\eta) + \mu_p \sigma_p \phi(\eta), \quad (17)$$

where $\eta = \mu_p/\sigma_p$, $\phi(\cdot)$ is the density function of the standard normal distribution and $\Phi(\cdot)$ is the cumulative standard normal distribution function.

In addition to Lemma 4, we also need the following lemma to solve the minimum truncated second moment portfolio problem.

Lemma 5. *Let $g(u) = u + \phi(u)/\Phi(u)$. We have (1) $g(u)$ is a positive and strictly increasing function of u , and (2) $\lim_{u \rightarrow -\infty} g(u) = 0$ and $\lim_{u \rightarrow \infty} g(u) = \infty$. It follows that $g(u) = c$ has a unique solution for $c > 0$.*

With the results in Lemma 4 and Lemma 5, we present the explicit solution to the constrained HJ-bound under the normality assumption in Proposition 1.

Proposition 1. *Suppose that $R \sim N(\mu, V)$. Let η^* be the unique solution of*

$$g(\eta) = \frac{1}{\theta_0}, \quad (18)$$

where θ_0 is defined in (7). Then, the portfolio with minimum truncated second moment has the following weights in the N risky assets

$$w^* = -\frac{R_0}{\theta_0(\eta^* + \theta_0)} V^{-1} (\mu - R_0 \mathbf{1}_N). \quad (19)$$

In addition, the constrained HJ-bound is given by

$$\text{Var}[m_c] = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)} - \frac{1}{R_0^2} = \frac{\sigma_0(\eta^* \mu_m + \sigma_0)}{\Phi(\eta^*)} - \mu_m^2 \equiv \sigma_c^2, \quad (20)$$

where $\mu_m = 1/R_0$ and σ_0^2 is the unconstrained HJ-bound defined in (10).

By examining (19), we can see that under the normality assumption, the minimum truncated second moment portfolio, just like the minimum second moment portfolio in (6), is a linear combination of the tangency portfolio and the risk-free asset. From (18), we know

$$\frac{1}{\theta_0} = \eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} > \eta^*, \quad (21)$$

and we have

$$0 < \theta_0(\eta^* + \theta_0) < 1 + \theta_0^2. \quad (22)$$

Consequently, the minimum truncated second moment portfolio in (19) involves short-selling more of the tangency portfolio than the minimum second moment portfolio in (6). With the nonnegativity constraint, σ_c^2 is naturally greater than σ_0^2 . From (18), we can see that η^* is a monotonically decreasing function of θ_0 . Therefore, for a given R_0 , the constrained HJ-bound σ_c^2 is uniquely determined by θ_0 , the Sharpe ratio of the tangency portfolio. The following lemma further shows that σ_c^2 is a monotonic increasing function of θ_0 . In addition, it shows that the difference between the constrained and unconstrained HJ-bounds is also a monotonically increasing function of θ_0 .

Lemma 6. *For a given R_0 , σ_c^2 is a strictly increasing function of θ_0 . In addition, let*

$$h(\theta_0) = \sigma_c^2 - \sigma_0^2, \quad \tilde{h}(\theta_0) = \frac{\sigma_c^2}{\sigma_0^2} - 1. \quad (23)$$

We have (1) $h(\theta_0)$ and $\tilde{h}(\theta_0)$ are positive and strictly increasing function of θ_0 . (2) $\lim_{\theta_0 \rightarrow 0} h(\theta_0) = \lim_{\theta_0 \rightarrow 0} \tilde{h}(\theta_0) = 0$, and (3) $\lim_{\theta_0 \rightarrow \infty} h(\theta_0) = \lim_{\theta_0 \rightarrow \infty} \tilde{h}(\theta_0) = \infty$.

Lemma 6 tells us that when the Sharpe ratio of the tangency portfolio is small, we expect the constrained and unconstrained HJ-bounds to be very close to each other. It is only when the tangency portfolio has a large Sharpe ratio that we can expect some meaningful difference between the two HJ-bounds.

With Proposition 1, we can now solve for the minimum truncated second moment portfolio q^* . This allows us to better understand the relation between the truncated Sharpe ratio of q^* and the constrained HJ-bound as given in Lemma 3. In Figure 4, we provide a graphical illustration of this relation in the space of (σ_p^+, μ_p^+) under the assumption that returns have a multivariate normal distribution. When the risk-free rate is R_0 , the two curves emanating from the point $(0, R_0)$ represent the minimum/maximum truncated variance frontier of the risk-free asset and the N

risky assets. It can be readily shown that just like the mean-variance frontier, the truncated mean-variance frontier is also a linear combination of the risk-free asset and the tangency portfolio, so that we can easily trace out the frontier by altering the weight in the risk-free asset. However, there are two major differences between the mean-variance frontier and the truncated mean-variance frontier. The first one is that the truncated mean-variance frontier is not represented by two straight lines as in the case of the mean-variance frontier. When the weight of the risk-free asset in a portfolio is close to one, the gross return on the portfolio has very small probability of assuming a negative value, and the truncated mean-variance frontier is almost identical to the mean-variance frontier when it is near the point $(0, R_0)$. When the portfolio has significant positive or negative weights in the tangency portfolio, then the truncated mean and standard deviation of the portfolio return can differ significantly from the mean and standard deviation. The second difference is that the truncated mean of the portfolio return has a lower bound, so the lower curve does not continue to go down but instead turns around after reaching a minimum.⁷ Once it turns around, the curve becomes the maximum truncated variance frontier rather than the minimum truncated variance frontier. The turnaround occurs because the probability of getting both large positive and negative returns increases with more short-selling of the tangency portfolio. Beyond a certain point, more short-selling of the tangency portfolio can actually increase the truncated expected return on the portfolio, since the negative returns are dropped in the calculation of the truncated mean.

Since $E[(R_p^+)^2] = (\mu_p^+)^2 + (\sigma_p^+)^2$, the portfolio with minimum truncated second moment has the shortest distance from the origin. In order to locate the portfolio with minimum truncated second moment, we draw a circle with its center at the origin, and the location of the minimum truncated second moment portfolio q^* can be obtained from the point where the circle is tangent to the minimum truncated variance frontier of the risk-free and risky assets. Note that the truncated Sharpe ratio of q^* as well as σ_{q^*}/μ_{q^*} are both equal to θ_c . Suppose that we draw a solid line joining the origin and q^* and a horizontal dotted line at the level of $1/R_0 = \mu_m$. Then, the intersection point of these two lines has a horizontal distance of $\mu_m\theta_c = \sigma_c$ from the y -axis. Therefore, this distance gives us the constrained HJ-bound on the standard deviation of nonnegative admissible SDFs.

Figure 4 about here

⁷Let u^* be the solution of the equation $\phi(u)/\Phi(u) = \theta_0$. It can be shown that $\min_w \mu_p^+ = R_0\Phi(u^*)$.

2. Sample Hansen-Jagannathan Bounds

Since the population HJ-bounds are unobservable, we need to estimate them using realized returns. Suppose that we have a time series of gross returns on the N risky assets, R_t , $t = 1, \dots, T$. We can estimate the mean and the covariance matrix of R_t using the GMM estimators

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T R_t, \quad (24)$$

$$\hat{V} = \frac{1}{T} \sum_{t=1}^T (R_t - \hat{\mu})(R_t - \hat{\mu})'. \quad (25)$$

In order to estimate the unconstrained HJ-bound for a given value of $\mu_m = 1/R_0$, researchers typically use the sample counterpart of (10)

$$\hat{\sigma}_0^2 = \hat{a}\mu_m^2 - 2\hat{b}\mu_m + \hat{c}, \quad (26)$$

where \hat{a} , \hat{b} , and \hat{c} are the sample estimators of the three efficiency set constants

$$\hat{a} = \hat{\mu}'\hat{V}^{-1}\hat{\mu}, \quad \hat{b} = 1'_N\hat{V}^{-1}\hat{\mu}, \quad \hat{c} = 1'_N\hat{V}^{-1}1_N. \quad (27)$$

For the constrained HJ-bound $\sigma_c^2 = \text{Var}[m_c]$, Hansen, Heaton, and Luttmer (1995) suggest a sample estimator of σ_c^2 which can be obtained in two steps. We first estimate $E[R_{q^*}^{+2}]$, where q^* is the minimum truncated second moment portfolio. This can be accomplished by computing the sample counterpart of (12)

$$\hat{\lambda} = \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 1_N)]^2, \quad (28)$$

where $R_0 = 1/\mu_m$. From property (3) of Lemma 2, we know that $E[m_c^2] = 1/E[R_{q^*}^{+2}]$. Therefore, using $\hat{\lambda}$ as an estimator of $E[R_{q^*}^{+2}]$, we can estimate σ_c^2 using

$$\hat{\sigma}_c^2 = \frac{1}{\hat{\lambda}} - \mu_m^2. \quad (29)$$

We call $\hat{\sigma}_c^2$ the nonparametric estimator of σ_c^2 because it does not require any knowledge of the joint distribution of the returns.

Note that in computing $\hat{\sigma}_c^2$, we need to solve for w^* in (28) numerically. Without a good initial estimate, numerical minimization can be time consuming and can also give us a local minimum rather than a global minimum. Based on our experience, the sample counterpart of (19) often

provides a good initial estimate of w^* and leads to fast convergence of the numerical minimization problem.

When returns are multivariate normally distributed, we propose an estimator of σ_c^2 that is simpler to compute and is more efficient than $\hat{\sigma}_c^2$. Following Proposition 1, we let $\hat{\eta}^*$ be the unique solution to

$$g(\eta) = \frac{\mu_m}{\hat{\sigma}_0}, \quad (30)$$

where $\hat{\sigma}_0^2$ is the sample unconstrained HJ-bound defined in (26). Using $\hat{\eta}^*$, we compute the maximum likelihood estimator of σ_c^2 as

$$\tilde{\sigma}_c^2 = \frac{\hat{\sigma}_0(\hat{\eta}^* \mu_m + \hat{\sigma}_0)}{\Phi(\hat{\eta}^*)} - \mu_m^2. \quad (31)$$

Unlike the nonparametric estimator $\hat{\sigma}_c^2$ which requires solving an N -dimensional minimization problem, the maximum likelihood estimator $\tilde{\sigma}_c^2$ requires solving only one nonlinear equation. As a result, $\tilde{\sigma}_c^2$ is significantly easier to obtain than $\hat{\sigma}_c^2$. In addition, there is only one solution to (30), so that we do not need to worry about getting a local minimum.

When returns are normally distributed, $\tilde{\sigma}_c^2$ is asymptotically more efficient than $\hat{\sigma}_c^2$. The reason is that we only need to estimate μ and V to obtain $\tilde{\sigma}_c^2$. In contrast, $\hat{\sigma}_c^2$ requires us to estimate the joint distribution of the returns, and there could be substantial volatility in its estimates. When returns are close to but not exactly normally distributed, we may still prefer to use $\tilde{\sigma}_c^2$ instead of $\hat{\sigma}_c^2$ because the latter can be very volatile. The finite sample performance of these two estimators under normal and non-normal distributions will be studied in Section 3.

2.1 Asymptotic Distributions

Traditionally, statistical inferences on the HJ-bounds are based on the asymptotic distributions of the sample HJ-bounds. In this section, we briefly review the existing asymptotic results on the sample unconstrained and constrained HJ-bounds. We then specialize the asymptotic results to the case where returns are multivariate normally distributed, under which we can derive analytical expressions for the asymptotic variance of the sample unconstrained and constrained HJ-bounds. The results under normality will be used in the next section, where we investigate how well the asymptotic distributions of these estimators approximate their finite sample distributions.

Under the assumptions that returns are jointly stationary and ergodic, and their fourth moments exist, Hansen, Heaton, and Luttmer (1995, Proposition 2.1) provide the asymptotic distributions of $\hat{\sigma}_0^2$ and $\hat{\sigma}_c^2$. Define $m_{0,t}$ and $m_{c,t}$ as the realizations of m_0 and m_c at time t

$$m_{0,t} = \frac{1 - (\mu - R_0 1_N)' V^{-1} (R_t - \mu)}{R_0} = \mu_m - (\mu_m \mu - 1_N)' V^{-1} (R_t - \mu), \quad (32)$$

$$m_{c,t} = \frac{R_{q^*,t}^+}{\|R_{q^*,t}^+\|^2}, \quad (33)$$

where q^* is the minimum truncated second moment portfolio. With some simplifications, the asymptotic results of Hansen, Heaton, and Luttmer (1995) can be written as

$$\sqrt{T}(\hat{\sigma}_0^2 - \sigma_0^2) \overset{A}{\approx} N(0, v_0), \quad (34)$$

$$\sqrt{T}(\hat{\sigma}_c^2 - \sigma_c^2) \overset{A}{\approx} N(0, v_c), \quad (35)$$

where $v_0 = \sum_{j=-\infty}^{\infty} E[\phi_{0,t} \phi_{0,t+j}]$ and $v_c = \sum_{j=-\infty}^{\infty} E[\phi_{c,t} \phi_{c,t+j}]$, and

$$\phi_{0,t} = m_{0,t}^2 - \mu_m^2 - \sigma_0^2, \quad (36)$$

$$\phi_{c,t} = m_{c,t}^2 - \mu_m^2 - \sigma_c^2. \quad (37)$$

In order to obtain a consistent estimator of v_0 , we can replace $\phi_{0,t}$ by

$$\hat{\phi}_{0,t} = \hat{m}_{0,t}^2 - \mu_m^2 - \hat{\sigma}_0^2, \quad (38)$$

where

$$\hat{m}_{0,t} = \frac{1 - (\hat{\mu} - R_0 1_N)' \hat{V}^{-1} (R_t - \hat{\mu})}{R_0} = \mu_m - (\mu_m \hat{\mu} - 1_N)' \hat{V}^{-1} (R_t - \hat{\mu}) \quad (39)$$

is the sample estimate of the minimum-variance SDF. Similarly, we can construct a consistent estimator of v_c by replacing $\phi_{c,t}$ with

$$\hat{\phi}_{c,t} = \hat{m}_{c,t}^2 - \mu_m^2 - \hat{\sigma}_c^2, \quad (40)$$

where

$$\hat{m}_{c,t} = \frac{R_{\hat{q}^*,t}^+}{\frac{1}{T} \sum_{t=1}^T (R_{\hat{q}^*,t}^+)^2} = \frac{R_{\hat{q}^*,t}^+}{R_0 \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+}, \quad (41)$$

and $R_{\hat{q}^*,t}^+$ is the gross return on the sample minimum truncated second moment portfolio at time t .⁸

When returns are i.i.d. multivariate normally distributed, we can derive analytical expressions of

⁸The last equality in (41) follows from the fact that the first order condition of the minimization problem is $\frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ (R_t - R_0 1_N) = 0_N$, which implies that

$$\frac{1}{T} \sum_{t=1}^T (R_{\hat{q}^*,t}^+)^2 = \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ R_{\hat{q}^*,t} = \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ [R_0 + w^{*'} (R_t - R_0 1_N)] = R_0 \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+. \quad (42)$$

the asymptotic variances of $\hat{\sigma}_0^2$ and $\hat{\sigma}_c^2$. In addition, we can also derive the asymptotic distribution of the maximum likelihood estimator of σ_c^2 . These asymptotic results are summarized in the following lemma.

Lemma 7. *Suppose that R_t , $t = 1, \dots, T$, are i.i.d. multivariate normally distributed. Then, the asymptotic distributions of $\hat{\sigma}_0^2$, $\hat{\sigma}_c^2$, and $\tilde{\sigma}_c^2$ are given by*

$$\sqrt{T}(\hat{\sigma}_0^2 - \sigma_0^2) \stackrel{A}{\approx} N\left(0, \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4}\right), \quad (43)$$

$$\sqrt{T}(\hat{\sigma}_c^2 - \sigma_c^2) \stackrel{A}{\approx} N\left(0, \frac{\theta_0^3[\theta_0(\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5)]}{R_0^4\Phi(\eta^*)^3} - \frac{\theta_0^2(\eta^* + \theta_0)^2}{R_0^4\Phi(\eta^*)^2}\right), \quad (44)$$

$$\sqrt{T}(\tilde{\sigma}_c^2 - \sigma_c^2) \stackrel{A}{\approx} N\left(0, \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4\Phi(\eta^*)^2}\right), \quad (45)$$

where θ_0^2 is defined in (7) and η^* is defined in Proposition 1. In addition, $\text{Avar}[\hat{\sigma}_c^2] \geq \text{Avar}[\tilde{\sigma}_c^2]$.

Since η^* is uniquely determined by θ_0 (the absolute value of the Sharpe ratio of the tangency portfolio of the risky assets), the asymptotic distributions of $\hat{\sigma}_0^2$, $\hat{\sigma}_c^2$, and $\tilde{\sigma}_c^2$ in Lemma 7 only depend on R_0 and θ_0 . In particular, these asymptotic distributions do not depend on N , the number of risky assets. This is in sharp contrast with our results in the next section which show that N plays a crucial role in determining the finite sample distributions of $\hat{\sigma}_0^2$, $\hat{\sigma}_c^2$, and $\tilde{\sigma}_c^2$.

3. Finite Sample Distributions of the Sample Hansen-Jagannathan Bounds

While the asymptotic distributions of the sample HJ-bounds are simple and easy to compute, they may not be reliable in finite samples. In this section, we present the finite sample distributions of $\hat{\sigma}_0^2$ and $\tilde{\sigma}_c^2$ under the normality assumption. For $\hat{\sigma}_c^2$, we cannot provide a simple expression of its finite sample distribution. Nevertheless, we are able to show that the finite sample distribution of $\tilde{\sigma}_c^2$ only depends on a single parameter. As a result, we are able to present a simple simulation approach to efficiently approximate the finite sample distribution of $\hat{\sigma}_c^2$.

Before analyzing the finite sample distributions of the sample HJ-bounds, we first present the finite sample distribution and moments of the squared sample Sharpe ratio of the tangency portfolio

$$\hat{\theta}_0^2 = (\hat{\mu} - R_0\mathbf{1}_N)' \hat{V}^{-1} (\hat{\mu} - R_0\mathbf{1}_N) \quad (46)$$

in the following proposition. The distribution of $\hat{\theta}_0^2$ can be easily obtained by using Theorem 3.2.12 of Muirhead (1982). It can also be obtained as a special case of the distribution of the Gibbons, Ross, and Shanken (1989) test when the number of factor mimicking portfolios is equal to zero.

Proposition 2. *Under the i.i.d. multivariate normality assumption on R_t , the distribution of the sample squared Sharpe ratio of the tangency portfolio is proportional to a noncentral F -distribution*

$$\hat{\theta}_0^2 \sim \left(\frac{N}{T-N} \right) F_{N, T-N}(T\theta_0^2), \quad (47)$$

where $F_{m,n}(\delta)$ denotes a noncentral F random variable with m and n degrees of freedom, and noncentrality parameter δ . The r -th moment of $\hat{\theta}_0$ exists if and only if $-N < r < T - N$. When $-N < r < T - N$, we have

$$E[\hat{\theta}_0^r] = \frac{\Gamma\left(\frac{N+r}{2}\right)\Gamma\left(\frac{T-N-r}{2}\right)}{\Gamma\left(\frac{N}{2}\right)\Gamma\left(\frac{T-N}{2}\right)} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right), \quad (48)$$

where $\Gamma(x)$ is the gamma function and ${}_1F_1(a; b; x)$ is the confluent hypergeometric function. When $r/2$ is a nonnegative integer, we have

$$E[\hat{\theta}_0^r] = \frac{\left(\frac{N}{2}\right)_{r/2}}{\left(\frac{T-N-r}{2}\right)_{r/2}} \sum_{i=0}^{r/2} \frac{\binom{r/2}{i} \left(\frac{T\theta_0^2}{2}\right)^i}{\left(\frac{N}{2}\right)_i}, \quad (49)$$

where $(n)_r = n(n+1)\cdots(n+r-1)$ and $(n)_0 = 1$.

3.1 Sample Unconstrained Hansen-Jagannathan Bound

Since $\hat{\sigma}_0^2 = \hat{\theta}_0^2/R_0^2$ is a linear transformation of $\hat{\theta}_0^2$, we can easily use the result in Proposition 2 to compute the finite sample distribution of $\hat{\sigma}_0^2$ as follows

$$P[\hat{\sigma}_0^2 < v] = P[\hat{\theta}_0^2 < R_0^2 v] = F_{N, T-N, T\theta_0^2} \left(\frac{(T-N)R_0^2 v}{N} \right), \quad (50)$$

where $F_{m,n,\delta}(x)$ is the noncentral F cumulative distribution function with m and n degrees of freedom, and noncentrality parameter δ .

In Figure 5, we plot the exact distribution of $\hat{\sigma}_0$ as a function of T for some representative values of θ_0 (0.2 and 0.4) and N (5 and 25). In each case, we assume $R_0 = 1.005$ and plot the population value of σ_0 using a horizontal solid line.⁹ We then plot the 1st, 5th, 50th, 95th, and

⁹Note that θ_0 of 0.2 and 0.4 cover a reasonably wide range of Sharpe ratios observed in monthly data. In addition, an R_0 of 1.005 seems sensible since it corresponds to an annual net return on the risk-free asset of about 6%.

99th percentiles of $\hat{\sigma}_0$ as functions of T . By comparing the four graphs in Figure 5, we can obtain a good understanding of how the finite sample distribution of $\hat{\sigma}_0$ varies with θ_0 , N , and T . In general, we see that the $\hat{\sigma}_0$ is not symmetrically distributed around σ_0 . The distribution of $\hat{\sigma}_0$ gets tighter as T increases but even for T as large as 600, there is still substantial volatility in $\hat{\sigma}_0$. By comparing the two upper panels ($\theta_0 = 0.2$) with the two lower panels ($\theta_0 = 0.4$) in Figure 5, we see that $\hat{\sigma}_0$ is more volatile for higher θ_0 but that the distribution of $\hat{\sigma}_0$ is more symmetric for higher θ_0 . By comparing the two left panels ($N = 5$) with the two right panels ($N = 25$) in Figure 5, we can see that an increase in N significantly increases the volatility of $\hat{\sigma}_0$ and drives the distribution of $\hat{\sigma}_0$ further away from σ_0 . When $\theta_0 = 0.2$ and $N = 25$, we notice that even for $T = 600$, the 1st percentile of $\hat{\sigma}_0$ is higher than the true σ_0 , indicating that $\hat{\sigma}_0$ provides a very poor estimate of σ_0 .

Figure 5 about here

The plots in Figure 5 suggest that there can be a significant upward bias in the distribution of $\hat{\sigma}_0$ especially when N is large. The underlying reason is that the sample tangency portfolio represents the outcome of an optimization problem that uses the sample mean and covariance matrix of returns. Since the optimizer tends to put heavy weights on assets with high average returns (but not necessarily high expected returns), the Sharpe ratio of the sample tangency portfolio ($\hat{\theta}_0$) tends to be considerably higher than the true θ_0 , especially when the number of assets is large.

Using the exact moments of $\hat{\theta}_0$ in (49), we can obtain the exact mean and variance of the sample unconstrained HJ-bound $\hat{\sigma}_0^2$ as¹⁰

$$E[\hat{\sigma}_0^2] = \frac{N + T\theta_0^2}{(T - N - 2)R_0^2} \quad \text{if } T - N > 2, \quad (52)$$

$$\text{Var}[\hat{\sigma}_0^2] = \frac{2[(N + T\theta_0^2)^2 + (N + 2T\theta_0^2)(T - N - 2)]}{(T - N - 2)^2(T - N - 4)R_0^4} \quad \text{if } T - N > 4. \quad (53)$$

Therefore, the ratios of the asymptotic mean and variance to the exact mean and variance of $\hat{\sigma}_0^2$ are given by

$$\frac{\sigma_0^2}{E[\hat{\sigma}_0^2]} = \frac{(T - N - 2)\theta_0^2}{N + T\theta_0^2}, \quad (54)$$

¹⁰Using the mean of $\hat{\sigma}_0^2$, we can easily derive an unbiased estimator of σ_0^2 as suggested by Ferson and Siegel (2003)

$$\hat{\sigma}_{0u}^2 = \left(\frac{T - N - 2}{T} \right) \hat{\sigma}_0^2 - \frac{N}{T} \mu_m^2. \quad (51)$$

$$\frac{\text{Avar}[\hat{\sigma}_0^2]}{\text{Var}[\hat{\sigma}_0^2]} = \frac{(T - N - 2)^2(T - N - 4)\theta_0^2(2 + \theta_0^2)}{T[(N + T\theta_0^2)^2 + (N + 2T\theta_0^2)(T - N - 2)]}. \quad (55)$$

These two ratios are only functions of N , T , and θ_0 . It is easy to show that both ratios are less than one, indicating that the exact distribution of $\hat{\sigma}_0^2$ has a higher mean than σ_0^2 and that $\hat{\sigma}_0^2$ is more volatile than what is suggested by the variance of its asymptotic distribution. In Figure 6, we plot the ratios $\sigma_0^2/E[\hat{\sigma}_0^2]$ and $\text{Avar}[\hat{\sigma}_0^2]/\text{Var}[\hat{\sigma}_0^2]$ as functions of T for four different choices of number of assets ($N = 2, 5, 10, \text{ and } 25$). The top two panels present the plots for the case in which $\theta_0 = 0.2$ and the bottom two panels present the plots for the case in which $\theta_0 = 0.4$. As expected, Figure 6 shows that the asymptotic distribution of $\hat{\sigma}_0^2$ provides a better approximation to the finite sample distribution of $\hat{\sigma}_0^2$ when T increases. Comparing the upper panels with the lower panels in Figure 6, we also find that the asymptotic distribution of $\hat{\sigma}_0^2$ is more accurate when θ_0 is higher. Finally, the quality of the approximation of the asymptotic distribution significantly deteriorates with an increase in N . When $N = 25$, the asymptotic distribution of $\hat{\sigma}_0^2$ provides a poor approximation to the exact distribution of $\hat{\sigma}_0^2$ even for T as large as 600.

Figure 6 about here

3.2 Maximum Likelihood Estimator of the Constrained Hansen-Jagannathan Bound

While the maximum likelihood estimator of the constrained HJ-bound, $\tilde{\sigma}_c^2$, looks complicated, it is actually just a monotonic transformation of $\hat{\theta}_0^2$. To see this, we can rewrite $\tilde{\sigma}_c^2$ in (31) as

$$\tilde{\sigma}_c^2 = \frac{1}{R_0^2} \left[\frac{\hat{\theta}_0(\hat{\eta}^* + \hat{\theta}_0)}{\Phi(\hat{\eta}^*)} - 1 \right], \quad (56)$$

where $\hat{\eta}^*$ is the solution to $g(\eta) = 1/\hat{\theta}_0$. Using the same proof as in Lemma 6, we can show that $\tilde{\sigma}_c^2$ is a monotonic increasing function of $\hat{\theta}_0^2$. Denoting the monotonic relation between $\tilde{\sigma}_c^2$ and $\hat{\theta}_0^2$ by

$$\tilde{\sigma}_c^2 = f(\hat{\theta}_0^2), \quad \hat{\theta}_0^2 = f^{-1}(\tilde{\sigma}_c^2), \quad (57)$$

we can again use Proposition 2 and obtain the finite sample distribution of $\tilde{\sigma}_c^2$ as follows

$$P[\tilde{\sigma}_c^2 < v] = P[f^{-1}(\tilde{\sigma}_c^2) < f^{-1}(v)] = F_{N, T-N, T\theta_0^2} \left(\frac{(T - N)f^{-1}(v)}{N} \right). \quad (58)$$

The only difference is that in this case we need to numerically compute $f^{-1}(v)$. Since $f(x)$ is a monotonic increasing function of x , solving for $f^{-1}(v)$ is fast and numerically stable.

In Figure 7, we plot the exact distribution of $\tilde{\sigma}_c$ as a function of T for some representative values of θ_0 (0.2 and 0.4) and N (5 and 25). In each case, we assume that $R_0 = 1.005$ and plot the population value of σ_c using a horizontal solid line. We then plot the 1st, 5th, 50th, 95th, and 99th percentiles of $\tilde{\sigma}_c$ as functions of T . Since $\tilde{\sigma}_c$ is a monotonic transformation of $\hat{\sigma}_0$, the distributions of $\tilde{\sigma}_c$ in Figure 7 are quite similar to the distributions of $\hat{\sigma}_0$ in Figure 5. Using a proof similar to the one of Lemma 6, we can show that the difference between $\tilde{\sigma}_c$ and $\hat{\sigma}_0$ is large only when $\hat{\sigma}_0$ is large. As a result, the lower percentiles of $\tilde{\sigma}_c$ and $\hat{\sigma}_0$ are almost identical. However, the 95th and 99th percentiles of $\tilde{\sigma}_c$ are significantly larger than those of $\hat{\sigma}_0$, especially when N and θ_0 are large. Despite the difference in the right tails of the distributions of $\tilde{\sigma}_c$ and $\hat{\sigma}_0$, the general pattern that we observe in the distribution of $\hat{\sigma}_0$ continues to hold for the distribution of $\tilde{\sigma}_c$. Namely, the distribution of $\tilde{\sigma}_c$ has a significant positive bias, and this bias becomes more severe when N is large and T is small. Similar to the $\hat{\sigma}_0$ case, the percentage bias of $\tilde{\sigma}_c$ is larger for smaller θ_0 .

Figure 7 about here

In Figure 8, we plot the ratios $\sigma_c^2/E[\tilde{\sigma}_c^2]$ and $\text{Avar}[\tilde{\sigma}_c^2]/\text{Var}[\tilde{\sigma}_c^2]$ as functions of T for four different choices of number of assets ($N = 2, 5, 10,$ and 25). The top two panels present the plots for the case of $\theta_0 = 0.2$ and the bottom two panels are for the case of $\theta_0 = 0.4$. Although we can compute the exact distribution of $\tilde{\sigma}_c$, it is not easy to obtain simple expressions for the finite sample mean and variance of $\tilde{\sigma}_c^2$. Therefore, we use the sample mean and variance of 100,000 simulations of $\tilde{\sigma}_c^2$ to approximate $E[\tilde{\sigma}_c^2]$ and $\text{Var}[\tilde{\sigma}_c^2]$.

Figure 8 about here

The plots in Figure 8 are very similar to the plots in Figure 6. They show that the asymptotic distribution of $\tilde{\sigma}_c^2$ does not always provide a good approximation to the finite sample distribution of $\tilde{\sigma}_c^2$. The quality of the approximation improves with larger T , larger θ_0^2 , and smaller N . When $N = 25$, the asymptotic distribution of $\tilde{\sigma}_c^2$ is unreliable even for T as large as 600.

Since $\tilde{\sigma}_c^2$ can be a heavily biased estimator of σ_c^2 , it is desirable to obtain an approximate unbiased estimator of σ_c^2 . Let $\sigma_c^2 = f(\theta_0^2)$ and $\tilde{\sigma}_c^2 = f(\hat{\theta}_0^2)$. Using a first-order Taylor series expansion and the fact that $\partial\sigma_c^2/\partial\theta_0^2 = 1/[R_0^2\Phi(\eta^*)]$, we have

$$\tilde{\sigma}_c^2 = f(\hat{\theta}_0^2) \approx f(\theta_0^2) + f'(\theta_0^2)(\hat{\theta}_0^2 - \theta_0^2) = \sigma_c^2 + \frac{1}{R_0^2\Phi(\eta^*)}(\hat{\theta}_0^2 - \theta_0^2). \quad (59)$$

Taking the expectation of both sides, we obtain

$$E[\tilde{\sigma}_c^2] \approx \sigma_c^2 + \frac{1}{R_0^2 \Phi(\eta^*)} (E[\hat{\theta}_0^2] - \theta_0^2) = \sigma_c^2 + \frac{N + (N + 2)\theta_0^2}{(T - N - 2)R_0^2 \Phi(\eta^*)}. \quad (60)$$

Therefore, we can replace θ_0^2 and η^* by their sample counterparts and use

$$\tilde{\sigma}_c^2 - \frac{N + (N + 2)\hat{\theta}_0^2}{(T - N - 2)R_0^2 \Phi(\hat{\eta}^*)} \quad (61)$$

as an approximate unbiased estimator of σ_c^2 . However, $\hat{\theta}_0^2$ can be a heavily upward biased estimator of θ_0^2 , especially when N is large. As a result, the above estimator tends to over-adjust and it can be biased downward when N is large. Another problem with the above estimator is that it can be negative. To correct for these two problems, we propose the use of

$$\tilde{\sigma}_{cu}^2 = \max \left[0, \tilde{\sigma}_c^2 - \frac{N + (N + 2)\hat{\theta}_{0u}^2}{(T - N - 2)R_0^2 \Phi(\hat{\eta}_u^*)} \right] \quad (62)$$

as an adjusted estimator of σ_c^2 , where

$$\hat{\theta}_{0u}^2 = \max \left[0, \left(\frac{T - N - 2}{T} \right) \hat{\theta}_0^2 - \frac{N}{T} \right] \quad (63)$$

and $\hat{\eta}_u^*$ is a solution to the following equation

$$u + \frac{\phi(u)}{\Phi(u)} = \frac{1}{\hat{\theta}_{0u}}. \quad (64)$$

In Figure 9, we plot the exact distribution of $\tilde{\sigma}_{cu}$ as a function of T for some representative values of θ_0 (0.2 and 0.4) and N (5 and 25). In each case, we assume that $R_0 = 1.005$ and plot the population value of σ_c using a horizontal solid line. We then plot the 1st, 5th, 50th, 95th, and 99th percentiles of $\tilde{\sigma}_{cu}$ as functions of T . By comparing Figure 9 with Figure 7, we observe that $\tilde{\sigma}_{cu}$ is much better behaved than $\tilde{\sigma}_c$. For example, the four plots in Figure 9 show that the 50th percentile of $\tilde{\sigma}_{cu}$ is very close to σ_c , while the 50th percentile of the unadjusted $\tilde{\sigma}_c$ in Figure 7 is significantly higher than σ_c , even for T as large as 600. Overall, the distribution of $\tilde{\sigma}_{cu}$ tends to be more symmetric and less volatile than the distribution of $\tilde{\sigma}_c$.

Figure 9 about here

Using 100,000 simulations of $\tilde{\sigma}_{cu}^2$, we estimate $E[\tilde{\sigma}_{cu}^2]$ and $\text{Var}[\tilde{\sigma}_{cu}^2]$ and plot the ratios $\sigma_c^2/E[\tilde{\sigma}_{cu}^2]$ and $\text{Avar}[\tilde{\sigma}_{cu}^2]/\text{Var}[\tilde{\sigma}_{cu}^2]$ as functions of T for four different choices of number of assets ($N = 2, 5,$

10, and 25) in Figure 10. The top two panels present the plots for the case of $\theta_0 = 0.2$ and the bottom two panels are for the case of $\theta_0 = 0.4$. The plots of the ratios of the asymptotic mean to the finite mean in Figure 10 are substantially different from the ones in Figure 8. They show that $\tilde{\sigma}_{cu}^2$ is close to being an unbiased estimator of σ_c^2 , except for when $N = 25$ and T is very small. Although the asymptotic variance of $\tilde{\sigma}_{cu}^2$ is still substantially lower than the finite sample variance of $\tilde{\sigma}_{cu}^2$, the ratio is closer to one after the bias adjustment. This suggests that $\tilde{\sigma}_{cu}^2$ is not only effective in removing the bias of $\tilde{\sigma}_c^2$, but is also less volatile than $\tilde{\sigma}_c^2$. Therefore, we consider $\tilde{\sigma}_{cu}^2$ to be a superior estimator of the constrained HJ-bound than $\tilde{\sigma}_c^2$.

Figure 10 about here

3.3 Nonparametric Estimator of the Constrained Hansen-Jagannathan Bound

The last estimator that we consider is the nonparametric estimator of the constrained HJ-bound, $\hat{\sigma}_c^2$. Although we cannot obtain the exact distribution of $\hat{\sigma}_c^2$, we show in the following proposition that the distribution of $\hat{\sigma}_c^2$ only depends on θ_0 .

Proposition 3. *Under the i.i.d. multivariate normality assumption on R_t , the distribution of $\hat{\sigma}_c^2$ has the same distribution of*

$$\frac{1}{R_0^2} \left(\frac{1}{\tilde{\lambda}} - 1 \right), \quad (65)$$

where

$$\tilde{\lambda} = \min_{\tilde{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \tilde{w}' z_t]^2, \quad (66)$$

and $z_t \sim N([\theta_0, 0'_{N-1}]', I_N)$.

Proposition 3 provides us with an efficient way of simulating $\hat{\sigma}_c^2$. It suggests that one only needs to simulate z_t (whose distribution only depends on θ_0) for $t = 1, \dots, T$ to obtain a draw of $\hat{\sigma}_c^2$. Based on 100,000 simulations, we plot the exact distribution of $\hat{\sigma}_c$ as a function of T for some representative values of θ_0 (0.2 and 0.4) and N (5 and 25) in Figure 11. In each case, we assume that $R_0 = 1.005$ and plot the population value of σ_c using a horizontal solid line. We then plot the 1st, 5th, 50th, 95th, and 99th percentiles of $\hat{\sigma}_c$ as functions of T . By comparing Figure 7 with Figure 11, we can see that the nonparametric estimator $\hat{\sigma}_c$ tends to be more volatile than the maximum likelihood estimator $\tilde{\sigma}_c$. This is particularly the case when $N = 25$ and $T \leq 120$, where

we find that $\hat{\sigma}_c$ can often be very large. Note that $\hat{\sigma}_c$ is inversely related to $\hat{\lambda}$ in (28), which is a nonparametric estimator of $E[R_{q^*}^{+2}]$, where q^* is the minimum truncated second moment portfolio. However, when N is large and T is small, there is a high probability that we can find a portfolio that has negative gross returns in almost every period in the sample. When this occurs, we have $\hat{\lambda} \approx 0$ and this results in a very large value of $\hat{\sigma}_c$.

Figure 11 about here

The very fat right tail of $\hat{\sigma}_c^2$ renders the asymptotic distribution of $\hat{\sigma}_c^2$ grossly inappropriate for approximating the finite sample distribution of $\hat{\sigma}_c^2$, especially when N is large and T is small. In fact, we are able to establish that $P[\hat{\lambda} = 0] > 0$ so that $P[\hat{\sigma}_c^2 = \infty] > 0$, which in turn implies the nonexistence of moments for $\hat{\sigma}_c^2$.¹¹ To show this, we note that for any nonzero N -vector w_0 , we have

$$\begin{aligned}
P[\hat{\lambda} = 0] &= P\left[\min_w \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 \mathbf{1}_N)]^2 = 0\right] \\
&\geq P\left[\sum_{t=1}^T \max[0, R_0 + w_0'(R_t - R_0 \mathbf{1}_N)]^2 = 0\right] \\
&= \prod_{t=1}^T P[R_0 + w_0'(R_t - R_0 \mathbf{1}_N) \leq 0] \\
&= \Phi\left(-\frac{\mu_p}{\sigma_p}\right)^T, \tag{67}
\end{aligned}$$

where $\mu_p = R_0 + w_0'(\mu - R_0 \mathbf{1}_N)$ and $\sigma_p^2 = w_0' V w_0$. The second to last equality in the above equation follows from the independence property of R_t . While $P[\hat{\lambda} = 0]$ goes to zero as T increases, it remains nonzero for any finite T . As a result, the moments of $\hat{\sigma}_c^2$ do not exist and we can no longer study the finite sample mean and variance of $\hat{\sigma}_c^2$ as we do for $\tilde{\sigma}_c^2$. Note that the result that $\hat{\sigma}_c^2$ does not have any finite moment is quite general and is not limited to the normality case. For any joint distribution of R_t , if there exists a nonzero N -vector w_0 such that $P[R_0 + w_0'(R_t - R_0 \mathbf{1}_N) \leq 0] > 0$, then we have $P[\hat{\lambda} = 0] > 0$ and the moments of $\hat{\sigma}_c^2$ do not exist.

¹¹Burnside (1994) notes that $\hat{\lambda}$ can be equal to zero for some values of R_0 . In his simulations, he also finds that such an event occurs quite frequently. Our result is stronger in the sense that we show analytically that $P[\hat{\lambda} = 0] > 0$ for any value of R_0 .

3.4 Confidence Intervals

From the results in the previous subsections, we find that all the sample HJ-bounds are quite volatile and have a serious bias, especially when N is large and T is small. This problem is particularly serious in the case of the nonparametric estimator of the constrained HJ-bound because this estimator does not even have finite moments. Given the high volatility of the sample HJ-bounds, it is unwise to rely solely on the point estimator of the HJ-bound to make inferences. It would be ideal to have an interval estimator of the HJ-bounds to improve our understanding of where the true HJ-bounds may fall.

As it turns out, constructing confidence intervals for σ_0^2 and σ_c^2 is the same as constructing a confidence interval for θ_0^2 . Suppose that we can find a pair $(\underline{\theta}_0^2, \bar{\theta}_0^2)$ to form a $100(1 - \alpha)\%$ confidence interval for θ_0^2 , i.e.

$$P[\underline{\theta}_0^2 \leq \theta_0^2 \leq \bar{\theta}_0^2] = 1 - \alpha. \quad (68)$$

Then, using the fact that σ_0^2 and σ_c^2 are monotonic increasing transformations of θ_0^2 , we can obtain the confidence interval for σ_0^2 as $(\underline{\theta}_0^2/R_0^2, \bar{\theta}_0^2/R_0^2)$ and the confidence interval for σ_c^2 as $(f(\underline{\theta}_0^2), f(\bar{\theta}_0^2))$, where f is a function such that $f(\theta_0^2) = \sigma_c^2$.

From Proposition 2, we know that the finite sample distribution of $\hat{\theta}_0^2$ is proportional to a noncentral F -distribution with noncentrality parameter $T\theta_0^2$. Since the noncentral F -distribution is decreasing in its noncentrality parameter, we can use the statistical method (see, for example, Casella and Berger (1990, Section 9.2.3)) to construct a confidence interval for θ_0^2 .¹² Using this methodology, we first plot the $100(\alpha/2)$ and $100(1 - \alpha/2)$ percentiles of the distribution of $\hat{\theta}_0^2$ for different values of θ_0^2 . We then draw a horizontal line at the observed value of $\hat{\theta}_0^2$. This horizontal line will first intersect the $100(1 - \alpha/2)$ percentile line and then the $100(\alpha/2)$ percentile line of $\hat{\theta}_0^2$. The interval between these two intersection points gives us a $100(1 - \alpha)\%$ confidence interval for θ_0^2 . Mathematically, $\bar{\theta}_0^2$ and $\underline{\theta}_0^2$ are implicitly determined by the following equations

$$F_{N, T-N, \bar{\theta}_0^2}(x) = \frac{\alpha}{2}, \quad (69)$$

$$F_{N, T-N, \underline{\theta}_0^2}(x) = 1 - \frac{\alpha}{2}, \quad (70)$$

where $x = (T - N)\hat{\theta}_0^2/N$, $\bar{\delta} = T\bar{\sigma}_0^2/\mu_m^2$, and $\underline{\delta} = T\underline{\sigma}_0^2/\mu_m^2$. Note that since $F_{N, T-N, \delta}(x)$ is de-

¹²Lewellen, Nagel, and Shanken (2006) also use the statistical method to construct the confidence interval for the unexplained squared Sharpe ratio of an asset pricing model.

creasing in the noncentrality parameter δ , (70) will not have a nonnegative solution for $\underline{\theta}_0^2$ when $F_{N,T-N,0}(x) < 1 - \alpha/2$. In this case, we set $\underline{\theta}_0^2 = 0$. Similarly, if $F_{N,T-N,0}(x) < \alpha/2$, we cannot find a nonnegative solution for $\bar{\theta}_0^2$, and we set $\bar{\theta}_0^2 = 0$.¹³

3.5 Effects of Nonnormality and Conditional Heteroskedasticity

The distributional results on the sample HJ-bounds in this paper are derived under the strong assumption of i.i.d. multivariate normality. While we certainly do not think that returns are exactly i.i.d. normal, we view the normality assumption as a good working approximation for monthly and annual returns, which are used in most of the applications of the HJ-bounds. Nevertheless, we are interested in understanding how robust our results are to departures from the i.i.d. normality assumption. In particular, we are interested in return distributions with fat tails because the returns on financial assets often exhibit a leptokurtic behavior. We therefore study two alternative return distributions that exhibit leptokurtic behavior: (1) the case where returns are multivariate t distributed; and (2) the case where returns exhibit conditional heteroskedasticity of a GARCH type as in Bollerslev (1986).

3.5.1 Nonnormality

In this experiment, we study the case where returns are multivariate t distributed with five degrees of freedom. With the choice of five degrees of freedom, the t -distribution exhibits extreme fat tails and potentially presents a serious challenge for our finite sample results that are derived under the normality assumption. Since we cannot derive the finite sample distribution of $\hat{\sigma}_0^2$ and $\tilde{\sigma}_c^2$ under the multivariate t -distribution assumption, we have to rely on simulation.¹⁴ Using 100,000 simulations, we find that the distribution of $\hat{\theta}_0^2$ under the t -distribution assumption has a slightly fatter right tail than under the normality assumption. However, the noncentral F -distribution remains a very good approximation of the exact distribution of $\hat{\theta}_0^2$. To demonstrate this, we present in Figure 12 the coverage probabilities of the 90%, 95%, and 99% confidence intervals for σ_0^2 (which are exact under the normality assumption) when returns are multivariate t distributed with five degrees of freedom. We plot the probability of coverage of the three confidence intervals as a function of T for

¹³A Matlab program for computing the confidence intervals for σ_0^2 and σ_c^2 is available upon request.

¹⁴It can be shown that under the multivariate t -distribution assumption on R_t , the finite sample distribution of $\hat{\theta}_0^2$ (and hence $\hat{\sigma}_0^2$ and $\tilde{\sigma}_c^2$) only depends on θ_0 .

some representative values of θ_0 (0.2 and 0.4) and N (5 and 25). The plots show that the actual probabilities of coverage are quite close to the confidence levels. For the 99% confidence intervals, the coverage probability is almost exact. For the 90% and 95% confidence intervals, the coverage probabilities are almost exact for the case in which $\theta_0 = 0.2$ but off by about 1% to 2% for the case in which $\theta_0 = 0.4$. The reason why the coverage probabilities are slightly off when $\theta_0 = 0.4$ compared to the case where $\theta_0 = 0.2$ can be understood by noticing that \hat{V} is more volatile under the t -distribution assumption than under the normality assumption and that the sample squared Sharpe ratio of the tangency portfolio is given by $\hat{\theta}_0^2 = (\hat{\mu} - R_0 \mathbf{1}_N)' \hat{V}^{-1} (\hat{\mu} - R_0 \mathbf{1}_N)$. Therefore, the increased volatility of \hat{V} has a bigger impact on $\hat{\theta}_0^2$ when average excess returns are high (i.e., when θ_0 is high) than when average excess returns are low (i.e., when θ_0 is low). Consequently for larger θ_0 , $\hat{\theta}_0^2$ is more volatile under the t -distribution assumption than under the normality assumption and this leads to a decrease in the coverage probabilities of our confidence intervals.

Figure 12 about here

In Figure 13, we repeat the same exercise for the confidence intervals for σ_c^2 . One additional issue emerges when computing the probability of coverage of the confidence intervals for σ_c^2 : the population value of σ_c^2 under the multivariate t -distribution assumption is different from the one under the multivariate normality assumption. It can be shown that for the multivariate t -distribution with ν degrees of freedom, σ_c^2 is given by¹⁵

$$\sigma_c^2 = \frac{1}{R_0^2} \left[\frac{\theta_0(\theta_0 + \eta^*)}{\Phi_{\nu-2}(\eta^*)} - 1 \right], \quad (71)$$

where $\Phi_{\nu-2}(\cdot)$ is the cumulative distribution function of a standard t -distribution with $\nu - 2$ degrees of freedom, and η^* is the solution to the following equation

$$\frac{\eta \Phi_{\nu} \left(\left(\frac{\nu}{\nu-2} \right)^{\frac{1}{2}} \eta \right) + \phi_{\nu-2}(\eta)}{\Phi_{\nu-2}(\eta)} = \frac{1}{\theta_0}, \quad (72)$$

where $\phi_{\nu-2}(\cdot)$ is the density function of a standard t -distribution with $\nu - 2$ degrees of freedom. When $\theta_0 = 0.2$, σ_c is almost identical under the normality and the t -distribution assumptions. When $\theta_0 = 0.4$, $\sigma_c = 0.3983$ under the normality assumption but it increases to 0.4024 under the t -distribution assumption with five degrees of freedom.

¹⁵The proof of this result is available upon request.

Since there is only a small difference in the population value of σ_c^2 under the two distributional assumptions, Figure 13 shows that the probabilities of coverage of the confidence intervals for σ_c^2 are quite close to the confidence levels even when returns are multivariate t distributed. When T increases, we can see a small decrease in the probability of coverage for the case in which $\theta_0 = 0.4$ (due to the fact that our confidence intervals are designed to cover a slightly different σ_c^2), but the probability of coverage is still quite accurate for T as large as 600. For the 99% confidence intervals, the coverage probability is almost exact. Similar to the unconstrained case, the coverage probabilities of the 90% and 95% confidence intervals are almost exact for the case in which $\theta_0 = 0.2$ but off by about 1% to 2% for the case in which $\theta_0 = 0.4$.

Figure 13 about here

3.5.2 Conditional Heteroskedasticity

In this experiment, we introduce conditional heteroskedasticity in the return data generating process and investigate whether the noncentral F -distribution remains a good approximation of the exact distribution of $\hat{\theta}_0^2$. For modeling returns on financial assets, the GARCH(1,1) process proposed by Bollerslev (1986) has become a fairly popular choice in the literature. However, since we have multiple assets in our framework, we also need to make assumptions on the dynamics of the correlations of the returns on different pairs of assets. For simplicity, we use Bollerslev's (1990) constant correlation multivariate GARCH(1,1) model that assumes that these correlations are constant over time. Instead of assuming that excess returns follow a constant correlation multivariate GARCH(1,1), we assume that the following transformed excess returns follow a constant correlation multivariate GARCH(1,1)

$$z_t = P'V^{-\frac{1}{2}}r_t, \tag{73}$$

where r_t is an $N \times 1$ vector of excess returns, and P is an $N \times N$ orthonormal matrix with its first column equal to $V^{-\frac{1}{2}}(\mu - R_0 1_N)/\theta_0$. This assumption is made for convenience because it allows us to generate the time series of each element of the transformed returns independently using the following univariate GARCH(1,1) process

$$z_{it} = E[z_{it}] + \epsilon_t$$

$$\begin{aligned}\epsilon_t &\sim N(0, h_t) \\ h_t &= \omega_i + \alpha_i \epsilon_{t-1}^2 + \beta_i h_{t-1},\end{aligned}\tag{74}$$

where $\omega_i > 0$, $\alpha_i \geq 0$, $\beta_i \geq 0$, and $\alpha_i + \beta_i < 1$ for $i = 1, \dots, N$. Since the unconditional variance of z_{it} is equal to one, we set $\omega_i = 1 - \alpha_i - \beta_i$. In addition, we assume $E[z_{1t}] = \theta_0$, and $E[z_{it}] = 0$ for $i > 1$, so that the returns have the desired unconditional Sharpe ratio of θ_0 . In order to simulate z_t , we also need to choose the α_i and β_i parameters. We calibrate these parameters using the transformed excess returns on the 25 monthly Fama-French size and book-to-market portfolio returns over the post-World War II period (1946/1–2006/12).¹⁶ The average α_i and β_i parameter estimates across the 25 assets are 0.092 and 0.815, respectively. We use these estimated parameters to generate strings of simulated z_t . In Figure 14, we present the coverage probabilities of the 90%, 95%, and 99% confidence intervals for σ_0^2 when each of the transformed return series follows a GARCH(1,1) process. Using 100,000 simulations, we plot the probability of coverage of the three confidence intervals as a function of T for some representative values of θ_0 (0.2 and 0.4) and N (5 and 25). For the $N = 5$ case, we use only the first five elements of the simulated z_t , while for the $N = 25$ case we use all the elements of the simulated z_t . The plots show that the actual probabilities of coverage are quite close to the confidence levels. Similar to the t -distribution case, the coverage probability of the 99% confidence intervals is almost exact, while the coverage probabilities of the 90% and 95% confidence intervals are almost exact for the case in which $\theta_0 = 0.2$ but off by about 1% to 2% for the case in which $\theta_0 = 0.4$.

Figure 14 about here

In Figure 15, we repeat the same exercise for the confidence intervals for σ_c^2 . Since we cannot analytically derive the population value of σ_c^2 under the GARCH(1,1) assumption, we rely on a large-scale simulation. Similar to Ferson and Siegel (2003), we form artificial samples just like in the simulations, but with 1,000,000 observations. Then, using the nonparametric estimator of the constrained HJ-bound, we set the population values of σ_c^2 equal to the sample values in the artificial samples with 1,000,000 observations. Based on our parameter values, we have $\sigma_c = 0.1990$ for the $\theta_0 = 0.2$ case, and $\sigma_c = 0.3985$ for the $\theta_0 = 0.4$ case. These population values are very close to the ones that we would obtain if the returns were normally distributed. Since there is only a small

¹⁶We thank Kenneth French for making these data available through his website.

difference in the population value of σ_c^2 under the two distributional assumptions, Figure 15 shows that the probabilities of coverage of the confidence intervals for σ_c^2 are quite close to the confidence levels even when each element of the transformed returns z_t is GARCH(1,1) distributed. For the 99% confidence intervals, the coverage probability is almost exact. Similar to the unconstrained HJ-bound case, the coverage probabilities of the 90% and 95% confidence intervals are almost exact for the case in which $\theta_0 = 0.2$ but off by about 1% to 2% for the case in which $\theta_0 = 0.4$.

Figure 15 about here

In summary, the coverage probabilities of the confidence intervals proposed in this paper are quite accurate even when returns exhibit severe departures from the i.i.d. multivariate normality assumption. Hence, we expect our confidence intervals to have good coverage probabilities when using monthly data to estimate the unconstrained and constrained HJ-bounds. Our confidence interval analysis under normality would work even better when carried out on annual return data since the departures from the i.i.d. multivariate normality assumption are smaller for annual data. In cases where returns have extreme fat tails and when the population Sharpe ratio is rather large, the actual coverage probabilities of our confidence intervals for the HJ-bounds can be smaller than the stated confidence levels. In those cases, one may treat our confidence intervals for the HJ-bounds as conservative estimates of the uncertainty of the location of the population HJ-bounds.

4. Conclusions

In this paper, we provide a geometrical interpretation of the unconstrained and constrained HJ-bounds in the mean-variance and truncated mean-variance frontiers of returns, respectively. Under the multivariate normality assumption on returns, we present the finite sample distributions of the sample unconstrained and constrained HJ-bounds. In addition, we show that the moments of the traditional nonparametric estimator of the constrained HJ-bound do not exist in finite samples. To overcome this problem, we present a simpler and more reliable maximum likelihood estimator of the constrained HJ-bound. To correct for the finite sample bias in the maximum likelihood estimator, we also provide an approximate unbiased estimator of the constrained HJ-bound.

For typical number of assets and length of time series, the sample constrained and unconstrained

HJ-bounds are very volatile. To account for their sampling variability, we propose a simple method to construct confidence intervals for the unconstrained and constrained HJ-bounds. Given how easy it is to compute these confidence intervals, we believe it would be good practice to report both the point estimates and the confidence intervals for the HJ-bound. Using simulation experiments, we show that our confidence intervals have accurate probabilities of coverage even when the distribution of the returns has fat tails and exhibits conditional heteroskedasticity that follows a GARCH(1,1) process.

Appendix

Proof of Lemma 1: (1) Since $E[R] = \mu$, the expectation of the second term in (9) vanishes and we have $E[m_0] = 1/R_0$. (2) Using the fact that $E[R(R - \mu)'] = V$, we have $E[m_0R] = (\mu - (\mu - R_01_N))/R_0 = 1_N$. (3) $\|m_0\|^2 = E[R_{p^*}^2]/\|R_{p^*}\|^4 = 1/\|R_{p^*}\|^2$. (4) Using (3), we have $m_0/\|m_0\|^2 = (R_{p^*}/\|R_{p^*}\|^2)\|R_{p^*}\|^2 = R_{p^*}$. (5) Since R_{p^*} is the gross return on a portfolio, we have $E[mR_{p^*}] = 1$. Using (3), we obtain

$$E[mm_0] = \frac{E[mR_{p^*}]}{\|R_{p^*}\|^2} = \frac{1}{\|R_{p^*}\|^2} = E[m_0^2]. \quad (\text{A1})$$

This implies that

$$\text{Cov}[m, m_0] = E[mm_0] - E[m]E[m_0] = E[m_0^2] - E[m_0]^2 = \text{Var}[m_0], \quad (\text{A2})$$

$$0 \leq \text{Var}[m - m_0] = \text{Var}[m] - 2\text{Cov}[m, m_0] + \text{Var}[m_0] = \text{Var}[m] - \text{Var}[m_0]. \quad (\text{A3})$$

This completes the proof.

Proof of Lemma 2: (1) Differentiating $E[R_p^{+2}] = E[R_p^+R_p]$ with respect to w , we can easily show that the return on the minimum truncated second moment portfolio must satisfy the following first order condition

$$E[R_{q^*}^+(R - R_01_N)] = 0_N. \quad (\text{A4})$$

Using (A4), we obtain

$$E[R_{q^*}^{+2}] = E[R_{q^*}^+R_{q^*}] = E[R_{q^*}^+(R_0 + w^*(R - R_01_N))] = R_0E[R_{q^*}^+]. \quad (\text{A5})$$

It follows that $E[m_c] = E[R_{q^*}^+]/E[R_{q^*}^{+2}] = 1/R_0$. (2) Dividing the first order condition (A4) by $\|R_{q^*}^+\|^2$ and using $E[m_cR_0] = 1$, we obtain

$$E[m_c(R - R_01_N)] = 0_N \Rightarrow E[m_cR] = 1_N. \quad (\text{A6})$$

(3) $\|m_c\|^2 = E[R_{q^*}^{+2}]/\|R_{q^*}^+\|^4 = 1/\|R_{q^*}^+\|^2$. (4) Using (3), we have $m_c/\|m_c\|^2 = (R_{q^*}^+/\|R_{q^*}^+\|^2)\|R_{q^*}^+\|^2 = R_{q^*}^+$. (5) Since $m > 0$ and R_{q^*} is a portfolio return, we have $E[mR_{q^*}^+] \geq E[mR_{q^*}] = 1$. Using (3), we obtain

$$E[mm_c] = \frac{E[mR_{q^*}^+]}{\|R_{q^*}^+\|^2} \geq \frac{E[mR_{q^*}]}{\|R_{q^*}^+\|^2} = \frac{1}{\|R_{q^*}^+\|^2} = E[m_c^2]. \quad (\text{A7})$$

This implies that

$$\text{Cov}[m, m_c] = E[mm_c] - E[m]E[m_c] \geq E[m_c^2] - E[m_c]^2 = \text{Var}[m_c], \quad (\text{A8})$$

$$0 \leq \text{Var}[m - m_c] = \text{Var}[m] - 2\text{Cov}[m, m_c] + \text{Var}[m_c] \leq \text{Var}[m] - \text{Var}[m_c]. \quad (\text{A9})$$

This completes the proof.

Proof of Lemma 3: Using (A5), we can simplify the squared truncated Sharpe ratio of q^* to

$$\theta_c^2 = \frac{(E[R_{q^*}^+] - R_0)^2}{\text{Var}[R_{q^*}^+]} = \frac{(E[R_{q^*}^+] - R_0)^2}{R_0 E[R_{q^*}^+] - E[R_{q^*}^+]^2} = \frac{R_0 - E[R_{q^*}^+]}{E[R_{q^*}^+]}. \quad (\text{A10})$$

Similarly, we have

$$\frac{\text{Var}[R_{q^*}^+]}{E[R_{q^*}^+]^2} = \frac{E[R_{q^*}^{+2}] - E[R_{q^*}^+]^2}{E[R_{q^*}^+]^2} = \frac{R_0 E[R_{q^*}^+] - E[R_{q^*}^+]^2}{E[R_{q^*}^+]^2} = \frac{R_0 - E[R_{q^*}^+]}{E[R_{q^*}^+]}. \quad (\text{A11})$$

Finally, using properties 1 and 3 in Lemma 2 and (A5), we have

$$\frac{\sigma_c^2}{\mu_m^2} = \frac{E[m_c^2] - \mu_m^2}{\mu_m^2} = \frac{\frac{1}{E[R_{q^*}^+]^2} - \frac{1}{R_0^2}}{\frac{1}{R_0^2}} = \frac{R_0^2}{E[R_{q^*}^+]^2} - 1 = \frac{R_0^2}{R_0 E[R_{q^*}^+]} - 1 = \frac{R_0 - E[R_{q^*}^+]}{E[R_{q^*}^+]}. \quad (\text{A12})$$

Therefore, we have $\theta_c^2 = \text{Var}[R_{q^*}^+]/E[R_{q^*}^+]^2 = \sigma_c^2/\mu_m^2$. This completes the proof.

Proof of Lemma 4: We first present some truncated moments of a standard normal random variable that will be used throughout the Appendix. Suppose that $y \sim N(0, 1)$. Using integration by parts and the fact that $d\phi(y)/dy = -y\phi(y)$ and $\phi(-\eta) = \phi(\eta)$, we obtain

$$\int_{-\eta}^{\infty} y\phi(y)dy = -\phi(y)|_{-\eta}^{\infty} = \phi(\eta), \quad (\text{A13})$$

$$\int_{-\eta}^{\infty} y^2\phi(y)dy = -y\phi(y)|_{-\eta}^{\infty} + \int_{-\eta}^{\infty} \phi(y)dy = -\eta\phi(\eta) + \Phi(\eta), \quad (\text{A14})$$

$$\int_{-\eta}^{\infty} y^3\phi(y)dy = -y^2\phi(y)|_{-\eta}^{\infty} + 2 \int_{-\eta}^{\infty} y\phi(y)dy = (\eta^2 + 2)\phi(\eta), \quad (\text{A15})$$

$$\int_{-\eta}^{\infty} y^4\phi(y)dy = -y^3\phi(y)|_{-\eta}^{\infty} + 3 \int_{-\eta}^{\infty} y^2\phi(y)dy = -(\eta^3 + 3\eta)\phi(\eta) + 3\Phi(\eta). \quad (\text{A16})$$

Letting $y = (R_p - \mu_p)/\sigma_p$ and using (A13)–(A14), we obtain

$$\begin{aligned} E[R_p^+] &= \int_0^{\infty} R_p f(R_p) dR_p \\ &= \int_{-\eta}^{\infty} (\mu_p + \sigma_p y) \phi(y) dy \end{aligned}$$

$$= \mu_p \Phi(\eta) + \sigma_p \phi(\eta), \quad (\text{A17})$$

$$\begin{aligned} E[R_p^{+2}] &= \int_0^\infty R_p^2 f(R_p) dR_p \\ &= \int_{-\eta}^\infty (\mu_p + \sigma_p y)^2 \phi(y) dy \\ &= \mu_p^2 \Phi(\eta) + 2\mu_p \sigma_p \phi(\eta) + \sigma_p^2 [-\eta \phi(\eta) + \Phi(\eta)] \\ &= (\mu_p^2 + \sigma_p^2) \Phi(\eta) + \mu_p \sigma_p \phi(\eta). \end{aligned} \quad (\text{A18})$$

This completes the proof.

Proof of Lemma 5: Define

$$f(u) = \frac{\phi(u)}{\Phi(u)}. \quad (\text{A19})$$

For $u \geq 0$, it is obvious that $g(u) = u + f(u) > 0$. For $u < 0$, Gordon (1941) shows that

$$0 < g(u) \leq -\frac{1}{u}. \quad (\text{A20})$$

Therefore, we have $g(u) > 0$ for all u . $\lim_{u \rightarrow \infty} g(u) = \infty$ is obvious. Taking the limit of (A20), we have $\lim_{u \rightarrow -\infty} g(u) = 0$. Differentiating $g(u)$, we have

$$g'(u) = 1 + f'(u) = 1 - f(u) [u + f(u)] > 0 \quad (\text{A21})$$

by the inequality $1 - uf(u) - f(u)^2 > 0$ due to Birnbaum (1942) and Sampford (1953). This completes the proof.

Proof of Proposition 1: Using the following derivatives

$$\frac{\partial \mu_p}{\partial w} = \mu - R_0 1_N, \quad (\text{A22})$$

$$\frac{\partial \sigma_p^2}{\partial w} = 2Vw, \quad (\text{A23})$$

$$\frac{\partial \sigma_p}{\partial w} = \frac{1}{\sigma_p} Vw, \quad (\text{A24})$$

$$\frac{\partial \Phi(\eta)}{\partial w} = \phi(\eta) \left[\frac{(\mu - R_0 1_N) \sigma_p - \eta Vw}{\sigma_p^2} \right], \quad (\text{A25})$$

$$\frac{\partial \phi(\eta)}{\partial w} = -\eta \phi(\eta) \left[\frac{(\mu - R_0 1_N) \sigma_p - \eta Vw}{\sigma_p^2} \right], \quad (\text{A26})$$

and Lemma 4, we can show that

$$\frac{\partial E[R_p^{+2}]}{\partial w} = \frac{\partial [(\mu_p^2 + \sigma_p^2) \Phi(\eta) + \mu_p \sigma_p \phi(\eta)]}{\partial w} = 2\Phi(\eta) [Vw + \sigma_p g(\eta) (\mu - R_0 1_N)]. \quad (\text{A27})$$

Setting this equal to zero, we have the following first order condition for portfolio q^*

$$w^* = -\sigma_{q^*} g(\eta^*) V^{-1}(\mu - R_0 \mathbf{1}_N). \quad (\text{A28})$$

The variance of R_{q^*} is then given by the following identity

$$\sigma_{q^*}^2 = w^{*'} V w^* = \sigma_{q^*}^2 g(\eta^*)^2 (\mu - R_0 \mathbf{1}_N)' V^{-1} (\mu - R_0 \mathbf{1}_N) = \sigma_{q^*}^2 g(\eta^*)^2 \theta_0^2. \quad (\text{A29})$$

Since $g(\eta^*)$ and θ_0 are positive, we can take the square root of both sides and obtain

$$g(\eta^*) = \frac{1}{\theta_0}, \quad (\text{A30})$$

and w^* must satisfy

$$w^* = -\frac{\sigma_{q^*}}{\theta_0} V^{-1}(\mu - R_0 \mathbf{1}_N). \quad (\text{A31})$$

With this expression, the expected return on q^* is given by

$$\mu_{q^*} = R_0 + w^{*'}(\mu - R_0 \mathbf{1}_N) = R_0 - \frac{\sigma_{q^*}}{\theta_0} (\mu - R_0 \mathbf{1}_N)' V^{-1}(\mu - R_0 \mathbf{1}_N) = R_0 - \theta_0 \sigma_{q^*}. \quad (\text{A32})$$

Dividing both sides by σ_{q^*} and using the fact that $\eta^* = \mu_{q^*} / \sigma_{q^*}$, we obtain

$$\sigma_{q^*} = \frac{R_0}{\eta^* + \theta_0}. \quad (\text{A33})$$

Substituting this into (A31), the minimum truncated second moment portfolio is given by

$$w^* = -\frac{R_0}{\theta_0(\eta^* + \theta_0)} V^{-1}(\mu - R_0 \mathbf{1}_N). \quad (\text{A34})$$

Using (A5) and Lemma 3, we have

$$E[R_{q^*}^{+2}] = R_0 E[R_{q^*}^+] = R_0 [\mu_{q^*} \Phi(\eta^*) + \sigma_{q^*} \phi(\eta^*)] = R_0 \sigma_{q^*} \Phi(\eta^*) g(\eta^*) = \frac{R_0^2 \Phi(\eta^*)}{\theta_0(\eta^* + \theta_0)}, \quad (\text{A35})$$

where the last equality is obtained by using (A30) and (A33). From property (3) in Lemma 2, we have

$$E[m_c^2] = \frac{1}{E[R_{q^*}^{+2}]} = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)}. \quad (\text{A36})$$

Using property (1) in Lemma 2, we have $E[m_c] = 1/R_0 = \mu_m$. Using the fact that $\sigma_0 = \theta_0/R_0$, we obtain the expression for $\text{Var}[m_c]$. This completes the proof.

Proof of Lemma 6: We first derive $\partial\eta^*/\partial\theta_0$. Differentiating both sides of (18) with respect to η^* and using (A21) gives us

$$\begin{aligned} -\frac{1}{\theta_0^2} \frac{\partial\theta_0}{\partial\eta^*} &= 1 - \frac{\phi(\eta^*)}{\Phi(\eta^*)} \left[\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} \right] \\ \Rightarrow -\frac{1}{\theta_0^2} \frac{\partial\theta_0}{\partial\eta^*} &= 1 - \left(\frac{1}{\theta_0} - \eta^* \right) \frac{1}{\theta_0} \\ \Rightarrow \frac{\partial\eta^*}{\partial\theta_0} &= \frac{1}{1 - \theta_0(\eta^* + \theta_0)}. \end{aligned} \quad (\text{A37})$$

We then obtain the derivative of σ_c^2 with respect to θ_0 as

$$\frac{\partial\sigma_c^2}{\partial\theta_0} = \frac{\partial}{\partial\theta_0} \frac{\theta_0(\eta^* + \theta_0)}{R_0^2\Phi(\eta^*)} = \frac{\eta^* + \theta_0 \frac{\partial\eta^*}{\partial\theta_0} + 2\theta_0 - \theta_0(\eta^* + \theta_0) \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)}. \quad (\text{A38})$$

Using the fact that $\phi(\eta^*)/\Phi(\eta^*) = (1/\theta_0) - \eta^*$ and (A37), we can simplify the derivative to

$$\frac{\partial\sigma_c^2}{\partial\theta_0} = \frac{\eta^* + 2\theta_0 - \eta^*[1 - \theta_0(\eta^* + \theta_0)] \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)} = \frac{2\theta_0}{R_0^2\Phi(\eta^*)} > 0. \quad (\text{A39})$$

Using this derivative and the fact that $\sigma_0^2 = \theta_0^2/R_0^2$, we obtain

$$h'(\theta_0) = \frac{2\theta_0}{R_0^2\Phi(\eta^*)} - \frac{2\theta_0}{R_0^2} > 0 \quad (\text{A40})$$

because $0 < \Phi(\eta^*) < 1$. From Lemma 5, we know $\lim_{\theta_0 \rightarrow 0} \eta^* = \infty$ which implies $\lim_{\theta_0 \rightarrow 0} \phi(\eta^*) = 0$ and $\lim_{\theta_0 \rightarrow 0} \Phi(\eta^*) = 1$. Consequently, we have

$$\lim_{\theta_0 \rightarrow 0} \theta_0 \eta^* = \lim_{\theta_0 \rightarrow 0} \frac{\eta^*}{\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)}} = 1. \quad (\text{A41})$$

It follows that $\lim_{\theta_0 \rightarrow 0} \sigma_c^2 = \frac{1}{R_0^2} - \frac{1}{R_0^2} = 0$ and $\lim_{\theta_0 \rightarrow 0} h(\theta_0) = 0$. In order to show that $\lim_{\theta_0 \rightarrow \infty} h(\theta_0) = \infty$, it suffices to show that $h''(\theta_0) > 0$ because $h(\theta_0)$ is a strictly increasing function of θ_0 . The second derivative of $h(\theta_0)$ is given by

$$h''(\theta_0) = \frac{2 - 2\theta_0 \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)} - \frac{2}{R_0^2} = \frac{2[1 - \Phi(\eta^*)] - 2\theta_0 \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)}. \quad (\text{A42})$$

The first term is obviously positive. The second term is negative because $\partial\eta^*/\partial\theta_0 < 0$ using Lemma 5. It follows that $h''(\theta_0) > 0$.

Using the L'Hôpital's Rule, it is easy to show that $\lim_{\theta_0 \rightarrow 0} \tilde{h}(\theta_0) = \lim_{\theta_0 \rightarrow 0} \frac{1}{\Phi(\eta^*)} - 1 = 0$ and $\lim_{\theta_0 \rightarrow \infty} \tilde{h}(\theta_0) = \lim_{\theta_0 \rightarrow \infty} \frac{1}{\Phi(\eta^*)} - 1 = \infty$. The derivative of $\tilde{h}(\theta_0)$ is given by

$$\tilde{h}'(\theta_0) = \frac{1}{\sigma_0^4} \left[\frac{2\theta_0}{R_0^2\Phi(\eta^*)} \sigma_0^2 - \frac{2\theta_0}{R_0^2} \sigma_c^2 \right]$$

$$\begin{aligned}
&= \frac{2\theta_0}{\sigma_0^4 R_0^4} \left[\frac{\theta_0^2}{\Phi(\eta^*)} - \frac{\theta_0(\eta^* + \theta_0)}{\Phi(\eta^*)} + 1 \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} \left[\frac{\Phi(\eta^*)}{\theta_0} - \eta^* \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} \left[\Phi(\eta^*) \left(\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} \right) - \eta^* \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} [-\eta^* (1 - \Phi(\eta^*)) + \phi(\eta^*)]. \tag{A43}
\end{aligned}$$

Let $u = -\eta^*$ and using the fact that $1 - \Phi(-\eta^*) = \Phi(-\eta^*) = \Phi(u)$ and $\phi(\eta^*) = \phi(-\eta^*) = \phi(u)$, the term in the bracket can be written as

$$u\Phi(u) + \phi(u) = \Phi(u) \left(u + \frac{\phi(u)}{\Phi(u)} \right), \tag{A44}$$

which is greater than zero as shown in the proof of Lemma 5. It follows that $\tilde{h}(\theta_0)$ is a strictly increasing function of θ_0 . This completes the proof.

Proof of Lemma 7: Let $u_t = (\mu - R_0 1_N)' V^{-1} (R_t - \mu)$. Under the i.i.d. normality assumption, it is easy to verify that $E[u_t] = 0$, $E[u_t^2] = \theta_0^2$, $E[u_t^3] = 0$, and $E[u_t^4] = 3\theta_0^4$. Hence, we have

$$v_0 = E[\phi_{0,t}^2] = E \left[\left(\frac{(1 - u_t)^2}{R_0^2} - \frac{1 + \theta_0^2}{R_0^2} \right)^2 \right] = \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4}. \tag{A45}$$

Under the i.i.d. normality assumption, we have $v_c = E[\phi_{c,t}^2]$. In order to derive $E[\phi_{c,t}^2]$, we need explicit expressions of the following terms

$$E[m_{c,t}^2] = \frac{1}{\|R_{q^*,t}^+\|^2}, \tag{A46}$$

$$E[m_{c,t}^4] = \frac{E[(R_{q^*,t}^+)^4]}{\|R_{q^*,t}^+\|^8} = E[m_{c,t}^2]^4 E[(R_{q^*,t}^+)^4]. \tag{A47}$$

The first term is straightforward. Using (A36), we have

$$E[m_{c,t}^2] = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)}. \tag{A48}$$

The key is to derive $E[(R_{q^*,t}^+)^4]$. Suppose that $R_p \sim N(\mu_p, \sigma_p^2)$. Then, we have

$$\begin{aligned}
E[(R_p^+)^4] &= \int_{-\eta}^{\infty} (\mu_p + \sigma_p y)^4 \phi(y) dy \\
&= \mu_p^4 \Phi(\eta) + 4\mu_p^3 \sigma_p \phi(\eta) + 6\mu_p^2 \sigma_p^2 [-\eta \phi(\eta) + \Phi(\eta)] \\
&\quad + 4\mu_p \sigma_p^3 (\eta^2 + 2) \phi(\eta) + \sigma_p^4 [-(\eta^3 + 3\eta) \phi(\eta) + 3\Phi(\eta)] \\
&= [(\eta^4 + 6\eta^2 + 3)\Phi(\eta) + \eta(\eta^2 + 5)\phi(\eta)] \sigma_p^4 \tag{A49}
\end{aligned}$$

by applying (A13)–(A16). Using this result, we obtain

$$\begin{aligned} E[(R_{q^*,t}^+)^4] &= [(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)]\sigma_{q^*}^4 \\ &= \frac{[(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)]R_0^4}{(\eta^* + \theta_0)^4} \end{aligned} \quad (\text{A50})$$

by using the expression of σ_{q^*} in (A33). It follows that

$$E[m_{c,t}^4] = \frac{\theta_0^4[(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)]}{R_0^4\Phi(\eta^*)^4}. \quad (\text{A51})$$

Using the fact that $\phi(\eta^*)/\Phi(\eta^*) = 1/\theta_0 - \eta^*$, we can write

$$\begin{aligned} E[m_{c,t}^4] &= \frac{\theta_0^4 \left[(\eta^{*4} + 6\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5)(\frac{1}{\theta_0} - \eta^*) \right]}{R_0^4\Phi(\eta^*)^3} \\ &= \frac{\theta_0^4(\eta^{*2} + 3)}{R_0^4\Phi(\eta^*)^3} + \frac{\eta^*\theta_0^3(\eta^{*2} + 5)}{R_0^4\Phi(\eta^*)^3} \\ &= \frac{\theta_0^3[\theta_0(\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5)]}{R_0^4\Phi(\eta^*)^3}. \end{aligned} \quad (\text{A52})$$

It follows that

$$\begin{aligned} E[\phi_{c,t}^2] &= E[m_{c,t}^4] - 2(\sigma_c^2 + \mu_m^2)E[m_{c,t}^2] + (\sigma_c^2 + \mu_m^2)^2 \\ &= E[m_{c,t}^4] - E[m_{c,t}^2]^2 \\ &= \frac{\theta_0^3[\theta_0(\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5)]}{R_0^4\Phi(\eta^*)^3} - \frac{\theta_0^2(\eta^* + \theta_0)^2}{R_0^4\Phi(\eta^*)^2}. \end{aligned} \quad (\text{A53})$$

The asymptotic distribution of $\tilde{\sigma}_c^2$ can be easily obtained using the delta method. Note that both σ_0^2 and σ_c^2 are monotonic functions of θ_0 . Using the fact $\partial\sigma_0^2/\partial\theta_0 = 2\theta_0/R_0^2$ and $\partial\sigma_c^2/\partial\theta_0 = 2\theta_0/R_0^2\Phi(\eta^*)$ from (A39), we obtain $\partial\sigma_c^2/\partial\sigma_0^2 = 1/\Phi(\eta^*)$. Using the delta method, we obtain

$$\sqrt{T}(\tilde{\sigma}_c^2 - \sigma_c^2) \overset{A}{\rightsquigarrow} N\left(0, \frac{v_0}{\Phi(\eta^*)^2}\right). \quad (\text{A54})$$

Finally, $\text{Avar}[\hat{\sigma}_c^2] \geq \text{Avar}[\tilde{\sigma}_c^2]$ follows from the fact $\tilde{\sigma}_c^2$ is the maximum likelihood estimator of σ_c^2 and hence it is asymptotically most efficient. This completes the proof.

Proof of Proposition 2: Under the normality assumption, we have

$$\hat{\mu} \sim N(\mu, V/T), \quad (\text{A55})$$

$$T\hat{V} \sim W_N(T-1, V), \quad (\text{A56})$$

and they are independent of each other. Define $\tilde{\theta}_0^2 = (\hat{\mu} - R_0 \mathbf{1}_N)' V^{-1} (\hat{\mu} - R_0 \mathbf{1}_N)$. Using Theorem 3.2.12 of Muirhead (1982), we have

$$Y \equiv \frac{T\tilde{\theta}_0^2}{\hat{\theta}_0^2} \sim \chi_{T-N}^2, \quad (\text{A57})$$

and the ratio is independent of $\hat{\mu}$ and hence independent of $\tilde{\theta}_0^2$. From (A55), we have

$$\hat{\mu} - R_0 \mathbf{1}_N \sim N(\mu - R_0 \mathbf{1}_N, V/T). \quad (\text{A58})$$

Therefore,

$$X \equiv T\tilde{\theta}_0^2 \sim \chi_N^2(T\theta_0^2) \quad (\text{A59})$$

and is independent of Y . Together, we have

$$\hat{\theta}_0^2 = \frac{T\tilde{\theta}_0^2}{Y} = \frac{X}{Y} \sim \left(\frac{N}{T-N} \right) F_{N, T-N}(T\theta_0^2). \quad (\text{A60})$$

Using the independence between X and Y , the r -th moment of $\hat{\theta}_0$ is given by

$$E[\hat{\theta}_0^r] = E[X^{\frac{r}{2}}] E[Y^{-\frac{r}{2}}]. \quad (\text{A61})$$

Note that $E[X^{\frac{r}{2}}]$ exists if and only if $r > -N$ and $E[Y^{-\frac{r}{2}}]$ exists if and only if $r < T - N$. Using the moments of a noncentral chi-squared distribution (see, for example, Johnson, Kotz, and Balakrishnan (1995, p.450)), we have

$$E[X^{\frac{r}{2}}] = \frac{2^{\frac{r}{2}} \Gamma(\frac{N+r}{2})}{\Gamma(\frac{N}{2})} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right), \quad (\text{A62})$$

$$E[Y^{-\frac{r}{2}}] = \frac{2^{-\frac{r}{2}} \Gamma(\frac{T-N-r}{2})}{\Gamma(\frac{T-N}{2})}. \quad (\text{A63})$$

Combining these two expressions, we obtain

$$E[\hat{\theta}_0^r] = \frac{\Gamma(\frac{N+r}{2}) \Gamma(\frac{T-N-r}{2})}{\Gamma(\frac{N}{2}) \Gamma(\frac{T-N}{2})} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right). \quad (\text{A64})$$

When $r/2$ is a nonnegative integer, the confluent hypergeometric function can be simplified to

$${}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right) = \sum_{i=0}^{r/2} \frac{(-\frac{r}{2})_i}{(\frac{N}{2})_i} \frac{-\left(\frac{T\theta_0^2}{2}\right)^i}{i!} = \sum_{i=0}^{r/2} \frac{\binom{r/2}{i} \left(\frac{T\theta_0^2}{2}\right)^i}{\binom{N}{2}_i}. \quad (\text{A65})$$

This completes the proof.

Proof of Proposition 3: The nonparametric estimator of σ_c^2 is given by

$$\hat{\sigma}_c^2 = \frac{1}{R_0^2} \left(\frac{1}{\tilde{\lambda}} - 1 \right), \quad (\text{A66})$$

where $\tilde{\lambda}$ is defined as

$$\tilde{\lambda} = \frac{1}{R_0^2} \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 \mathbf{1}_N)]^2. \quad (\text{A67})$$

Defining $r_t = R_t - R_0 \mathbf{1}_N$ as the excess returns on the N risky assets at time t , we can rewrite $\tilde{\lambda}$ as

$$\tilde{\lambda} = \frac{1}{R_0^2} \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'r_t]^2 = \min_{\hat{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \hat{w}'r_t]^2, \quad (\text{A68})$$

where $\hat{w} = w/R_0$.

Letting $\tilde{\mu} = V^{-\frac{1}{2}}(\mu - R_0 \mathbf{1}_N)$, we have that $\tilde{\mu}'\tilde{\mu} = \theta_0^2$. Now, define P as an $N \times N$ orthonormal matrix with its first column equal to $\tilde{\mu}/(\tilde{\mu}'\tilde{\mu})^{\frac{1}{2}} = \tilde{\mu}/\theta_0$. Since $r_t \sim N(\mu - R_0 \mathbf{1}_N, V)$, we have that

$$z_t \equiv P'V^{-\frac{1}{2}}r_t \sim N([\theta_0, 0'_{N-1}]', I_N). \quad (\text{A69})$$

Having defined z_t , we can rewrite $\tilde{\lambda}$ as

$$\tilde{\lambda} = \min_{\tilde{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \tilde{w}'z_t]^2, \quad (\text{A70})$$

where $\tilde{w} = P'V^{\frac{1}{2}}\hat{w}$. This completes the proof.

References

- Backus, D. K., A. W. Gregory, and C. I. Telmer, 1993, "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance*, 48, 1887–1908.
- Balduzzi, P., and H. Kallal, 1997, "Risk Premia and Variance Bounds," *Journal of Finance*, 52, 1913–1949.
- Bekaert, G., and R. J. Hodrick, 1992, "Characterizing Predictable Components in Excess Returns on Equity and in Foreign Exchange Markets," *Journal of Finance*, 47, 467–510.
- Bekaert, G., and J. Liu, 2004, "Conditioning Information and Variance Bounds on Pricing Kernels," *Review of Financial Studies*, 17, 339–378.
- Birnbaum, Z. W., 1942, "An Inequality for Mills' Ratio," *Annals of Mathematical Statistics*, 13, 245–246.
- Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.
- Bollerslev, T., 1990, "Modelling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalized ARCH Model," *Review of Economics and Statistics*, 72, 498–505.
- Burnside, A. C., 1994, "Hansen-Jagannathan Bounds as Classical Tests of Asset Pricing Models," *Journal of Business and Economic Statistics*, 12, 57–79.
- Casella, G., and R. L. Berger, 1990, *Statistical Inference*, Duxbury Press, California.
- Cecchetti, S. G., P. Lam, and N. C. Mark, 1992, "Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns," *NBER* technical paper #124.
- Cecchetti, S. G., P. Lam, and N. C. Mark, 1994, "Testing Volatility Restrictions on Intertemporal Marginal Rates of Substitution Implied by Euler Equations and Asset Returns," *Journal of Finance*, 49, 123–152.
- Epstein, L. G., and S. E. Zin, 2001, "The Independence Axiom and Asset Returns," *Journal of Empirical Finance*, 8, 537–572.

- Ferson, W. E., and C. R. Harvey, 1992, "Seasonality and Consumption-based Asset Pricing," *Journal of Finance*, 47, 511–552.
- Ferson, W. E., and A. F. Siegel, 2003, "Stochastic Discount Factor Bounds with Conditioning Information," *Review of Financial Studies*, 16, 567–595.
- Gibbons, M. R., S. A. Ross, and J. Shanken, 1989, "A Test of the Efficiency of a Given Portfolio," *Econometrica*, 57, 1121–1152.
- Gordon, R. D., 1941, "Values of Mills' Ratio of Area to Bounding Ordinate of the Normal Probability Integral for Large Values of the Argument," *Annals of Mathematical Statistics*, 12, 364–366.
- Gregory, A. W., and G. W. Smith, 1992, "Sampling Variability in Hansen-Jagannathan Bounds," *Economics Letters*, 38, 263–267.
- Hagiwara, M., and M. A. Hecce, 1997, "Risk Aversion and Stock Price Sensitivity to Dividends," *American Economic Review*, 87, 738–745.
- Hansen, L. P., J. Heaton, and E. G. J. 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.
- He, H., and D. M. Modest, 1995, "Market Frictions and Consumption-based Asset Pricing," *Journal of Political Economy*, 103, 94–117.
- Heaton, J., 1995, "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," *Econometrica*, 63, 681–717.
- Johnson, N. L., S. Kotz, and N. Balakrishnan, 1995, *Continuous Univariate Distributions*, Vol. 2, Wiley, New York.
- Kan, R., and G. Zhou, 2006, "A New Variance Bound on the Stochastic Discount Factor," *Journal of Business*, 79, 941–961.

- Lewellen, J. W., S. Nagel, and J. Shanken, 2006, "A Skeptical Appraisal of Asset-pricing Tests," working paper, Dartmouth College.
- Mehra, R., and E. C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Muirhead, R. J., 1982, *Aspects of Multivariate Statistical Theory*, Wiley, New York.
- Otrok, C., B. Ravikumar, and C. H. Whiteman, 2002, "Evaluating Asset-pricing Models Using the Hansen-Jagannathan Bound: A Monte Carlo Investigation," *Journal of Applied Econometrics*, 17, 149–174.
- Sampford, M. R., 1953, "Some Inequalities on Mills' Ratio and Related Functions," *Annals of Mathematical Statistics*, 24, 130–132.
- Shiller, R. J., 1989, *Market Volatility*, MIT Press, Massachusetts.
- Snow, K. N., 1991, "Diagnosing Asset Pricing Models Using the Distribution of Asset Returns," *Journal of Finance*, 46, 955–984.

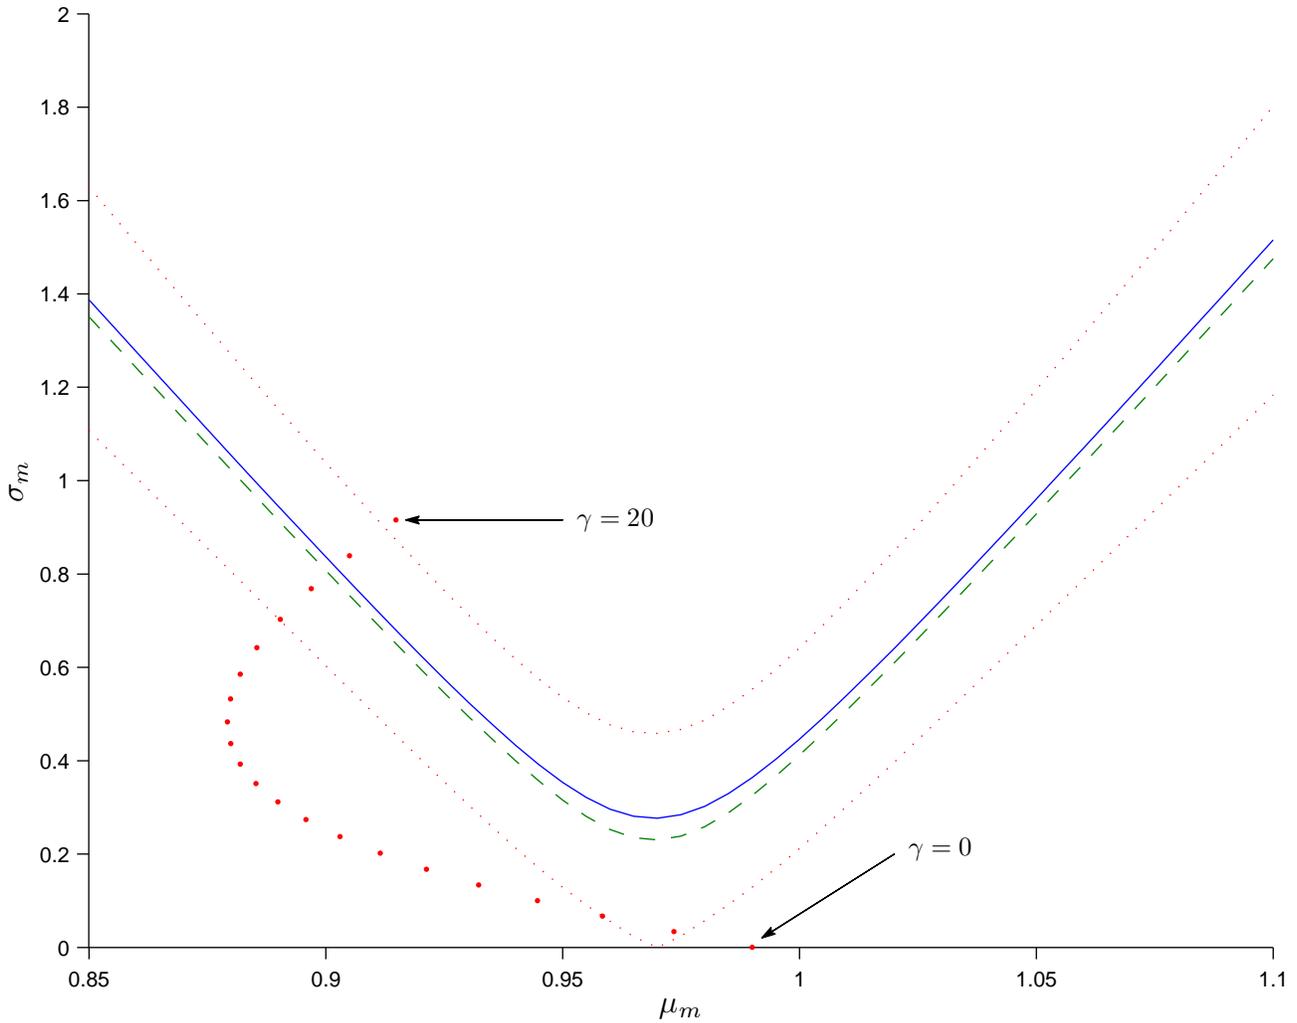


Figure 1

Confidence intervals for the unconstrained Hansen-Jagannathan bound. The figure presents the standard deviation and mean of the stochastic discount factor implied by a time-separable constant relative risk aversion utility function for different coefficients of relative risk aversion γ ranging between zero and 20 (in increments of one) and a subjective discount factor of 0.99. The figure also presents the sample unconstrained Hansen-Jagannathan bound (solid line), the bias-adjusted sample unconstrained Hansen-Jagannathan bound (dashed line), and the 95% confidence intervals (dotted lines) for the unconstrained Hansen-Jagannathan bound. The sample Hansen-Jagannathan bounds and the confidence intervals are computed using annual real returns on the S&P composite stock price index and commercial paper over the period 1891–1985.

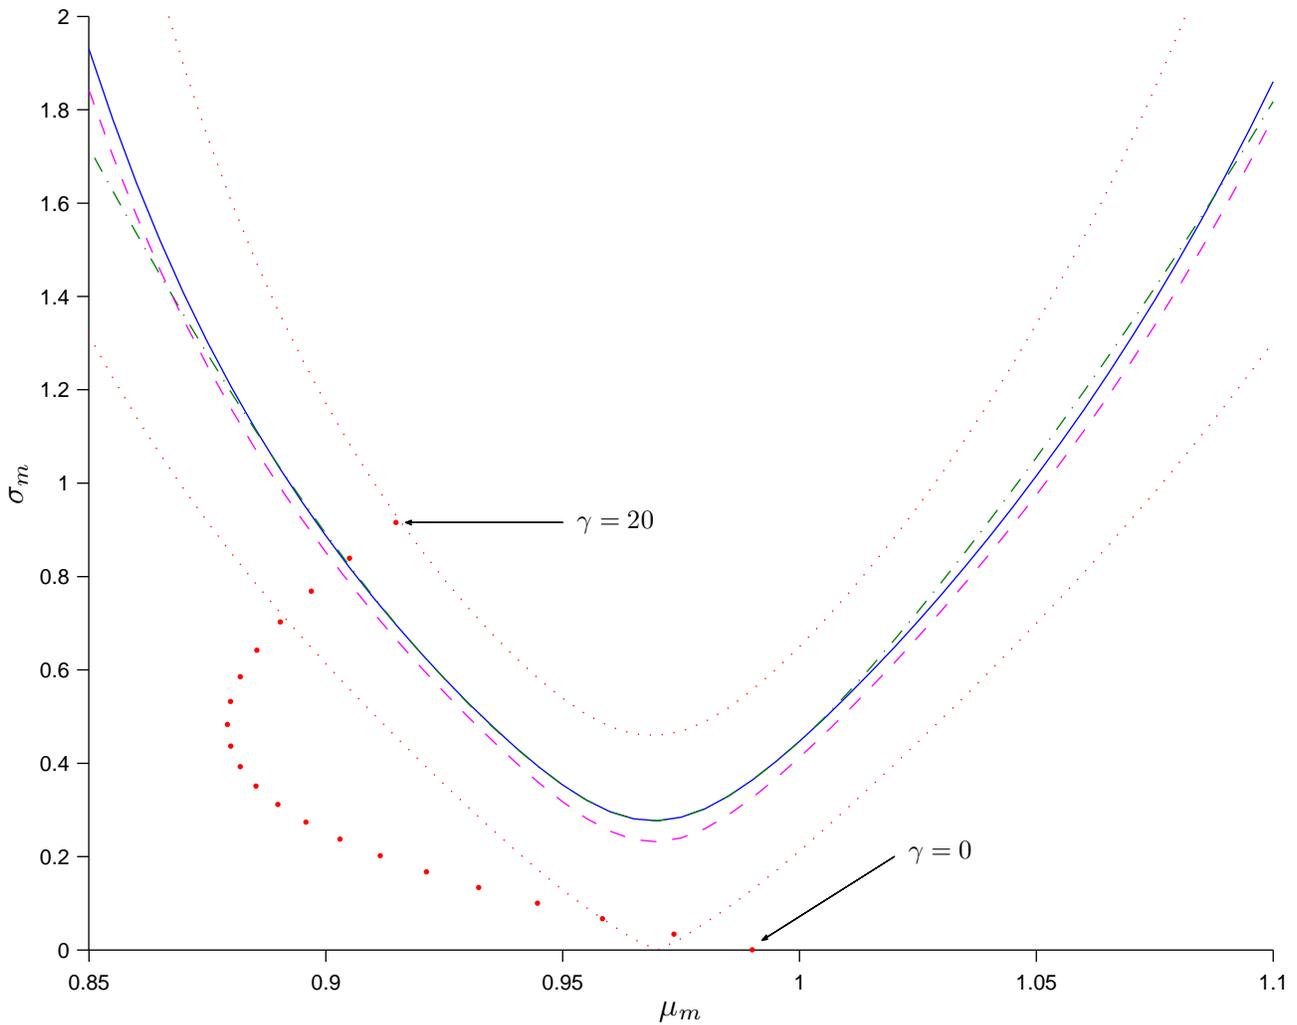


Figure 2

Confidence intervals for the constrained Hansen-Jagannathan bound. The figure presents the standard deviation and mean of the stochastic discount factor implied by a time-separable constant relative risk aversion utility function for different coefficients of relative risk aversion γ ranging between zero and 20 (in increments of one) and a subjective discount factor of 0.99. The figure also presents the maximum likelihood (solid line), the nonparametric (dashed-dotted line) and an approximate unbiased (dashed line) estimator of the constrained Hansen-Jagannathan bound as well as the 95% confidence intervals (dotted lines) for the constrained Hansen-Jagannathan bound. The sample constrained Hansen-Jagannathan bounds and the confidence intervals are computed using annual real returns on the S&P composite stock price index and commercial paper over the period 1891–1985.

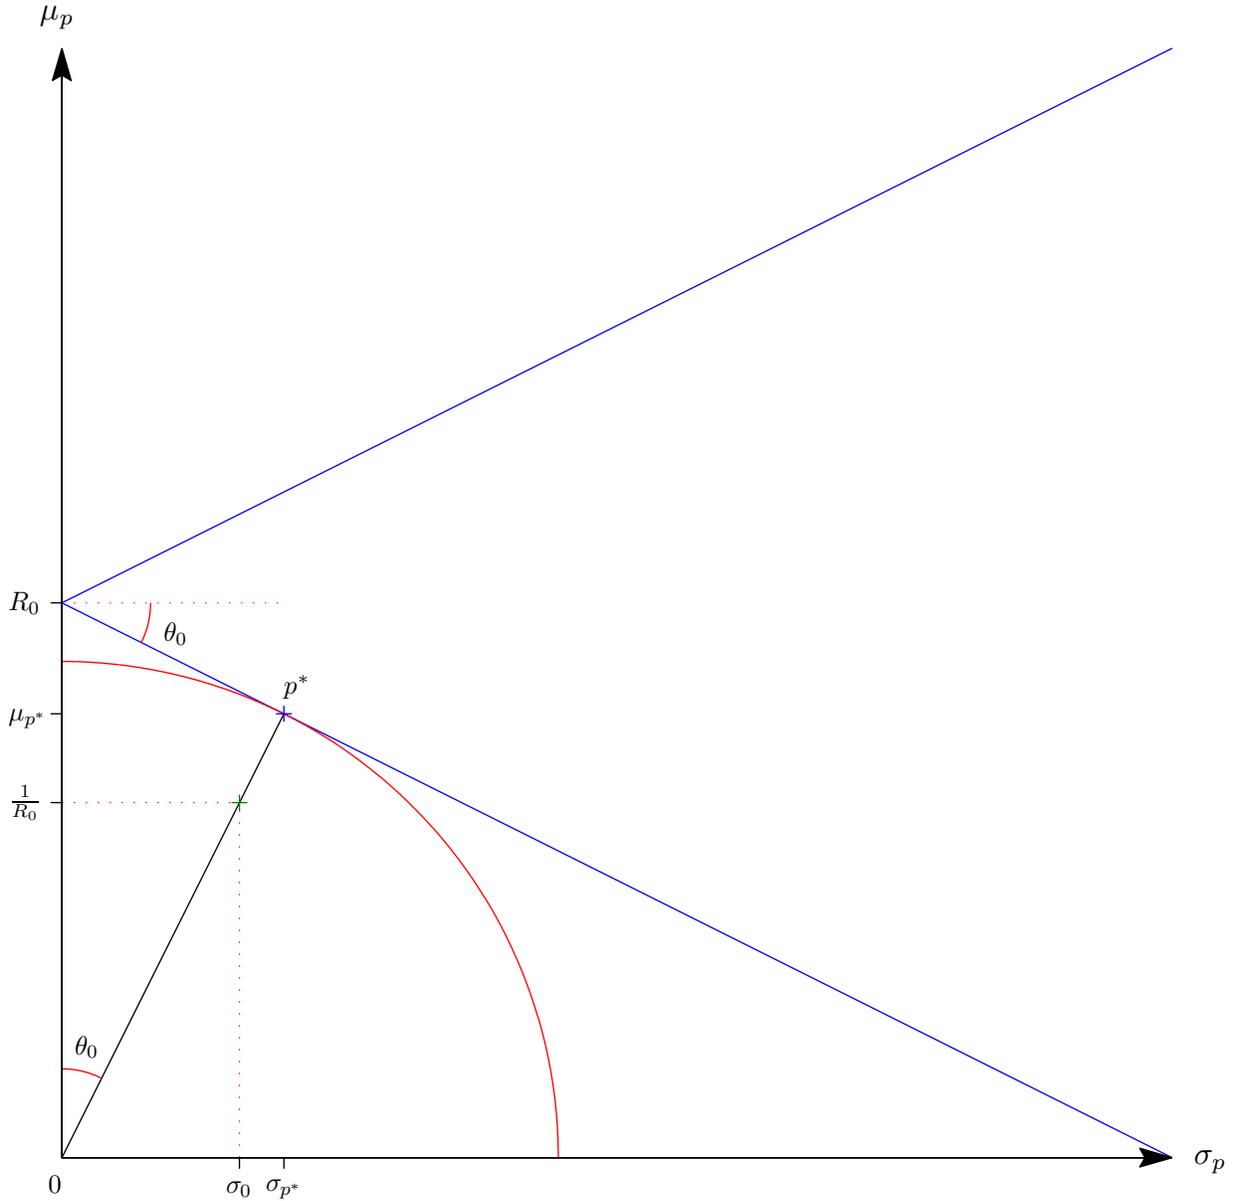


Figure 3

Graphical representation of the unconstrained Hansen-Jagannathan bound. The figure provides a geometrical interpretation of the unconstrained HJ-bound using the mean-variance frontier of portfolio returns. R_0 is the gross risk-free rate. The two solid lines emanating from the point $(0, R_0)$ represent the minimum-variance frontier of the risk-free and risky assets in the space of (σ_p, μ_p) . The portfolio with minimum second moment is represented by p^* and it is the portfolio that is closest to the origin. The absolute value of the Sharpe ratio of p^* as well as σ_{p^*}/μ_{p^*} are both equal to θ_0 . The horizontal distance between the point $(0, 1/R_0)$ and the solid line joining p^* and the origin is equal to σ_0 , where σ_0 is the unconstrained HJ-bound on the standard deviation of admissible SDFs with $E[m] = 1/R_0$.

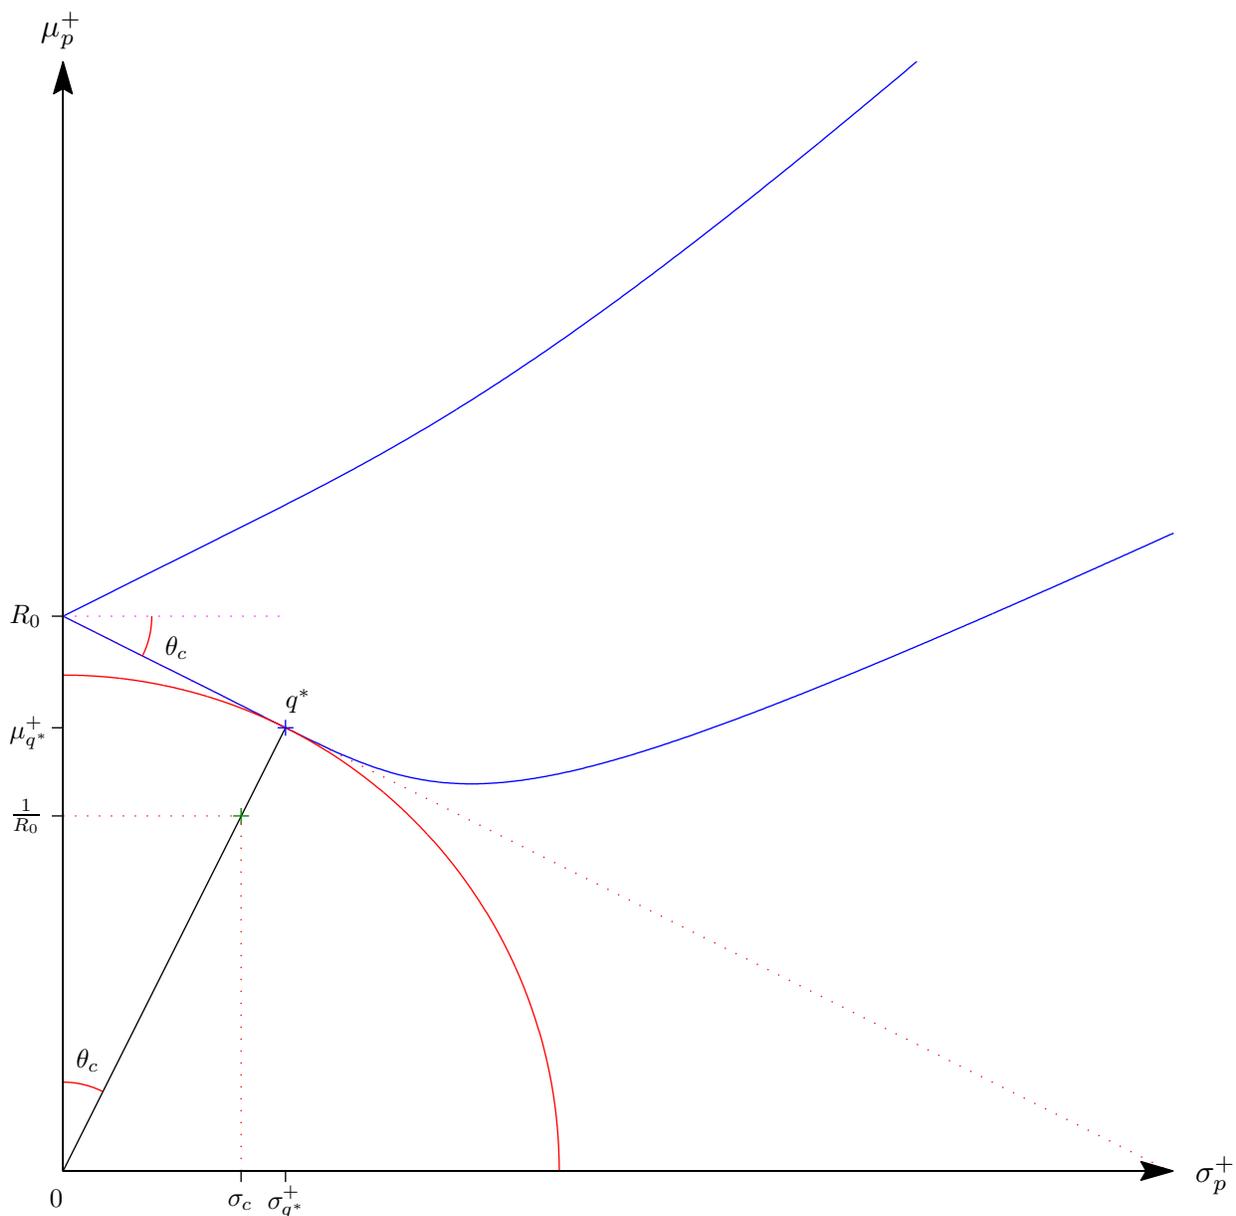


Figure 4

Graphical representation of the constrained Hansen-Jagannathan bound. The figure provides a geometrical interpretation of the constrained HJ-bound using the truncated mean-variance frontier of portfolio returns. R_0 is the gross risk-free rate. The two curves emanating from the point $(0, R_0)$ represent the minimum/maximum truncated variance frontier of the risk-free and risky assets in the space of (σ_p^+, μ_p^+) . The portfolio with minimum truncated second moment is represented by q^* and it is the portfolio that is closest to the origin. The absolute value of the truncated Sharpe ratio of q^* as well as $\sigma_{q^*}^+/\mu_{q^*}^+$ are both equal to θ_c . The horizontal distance between the point $(0, 1/R_0)$ and the solid line joining q^* and the origin is equal to σ_c , where σ_c is the HJ-bound on the standard deviation of nonnegative admissible SDFs with $E[m] = 1/R_0$.

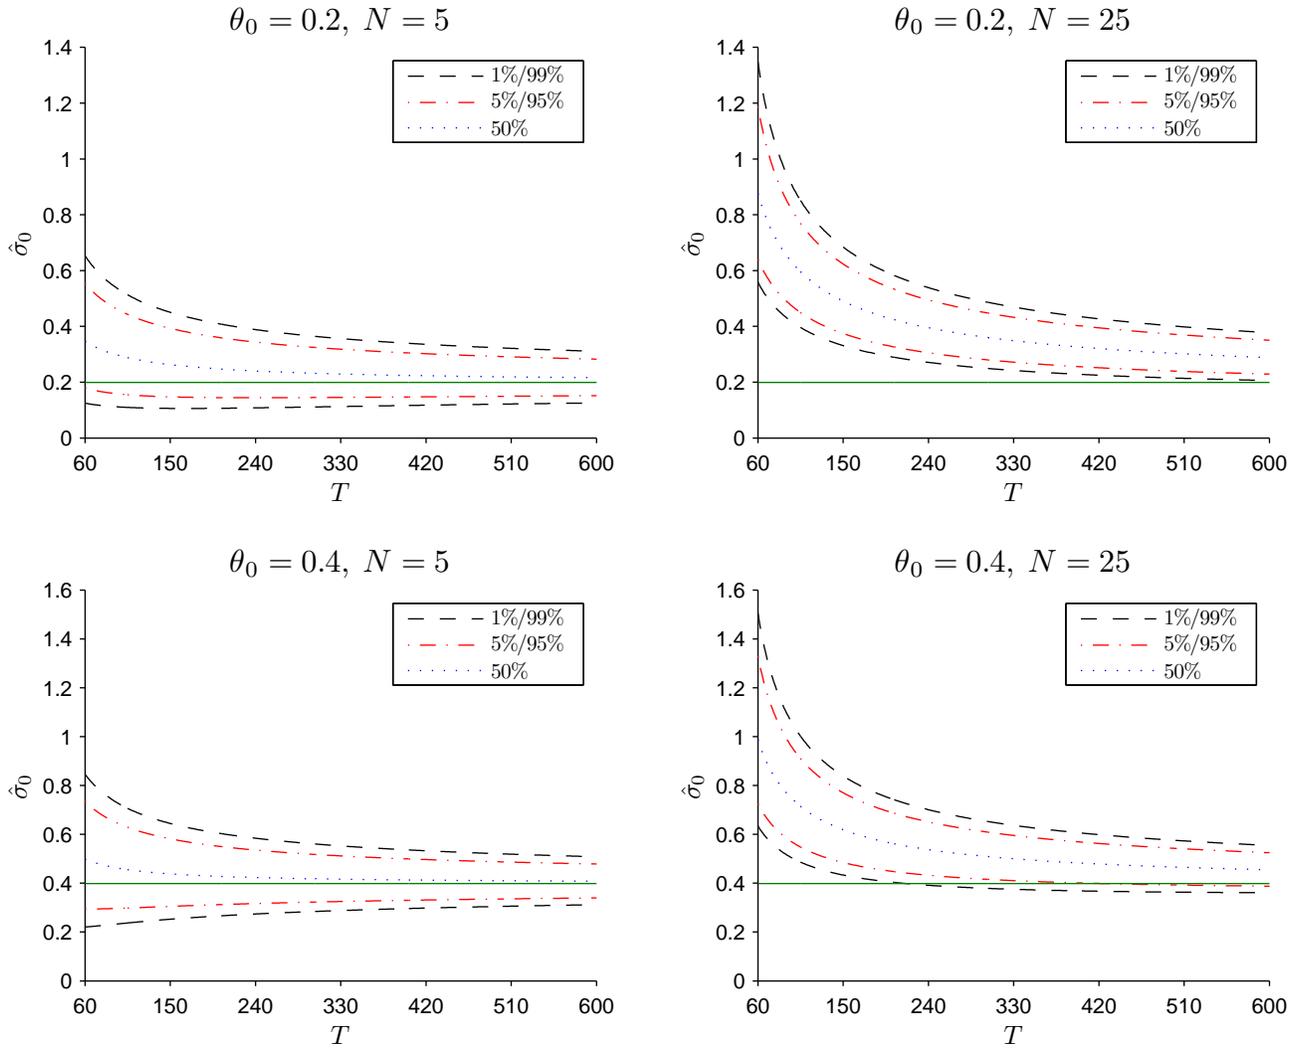


Figure 5

Exact distribution of the sample unconstrained Hansen-Jagannathan bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the sample unconstrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets. The gross risk-free rate is assumed to be 1.005 and the solid line in the figure represents the population value of the unconstrained HJ-bound.

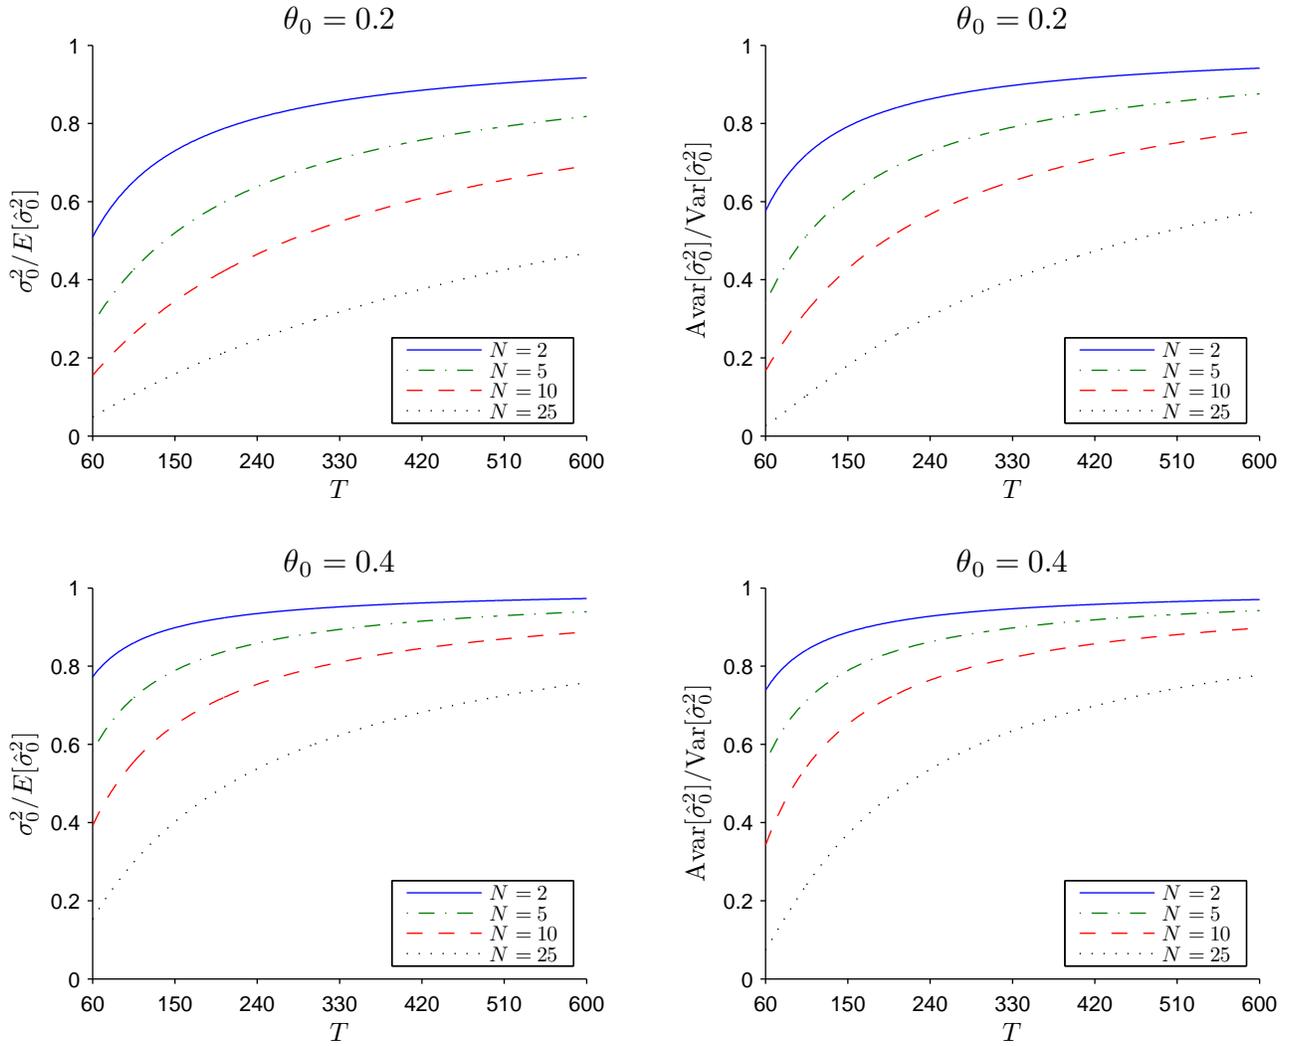


Figure 6

Ratio of asymptotic to exact moments of the sample unconstrained Hansen-Jagannathan bound. The figure presents the ratios of the asymptotic mean and variance to the exact mean and variance of the sample unconstrained HJ-bound ($\hat{\sigma}_0^2$) for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

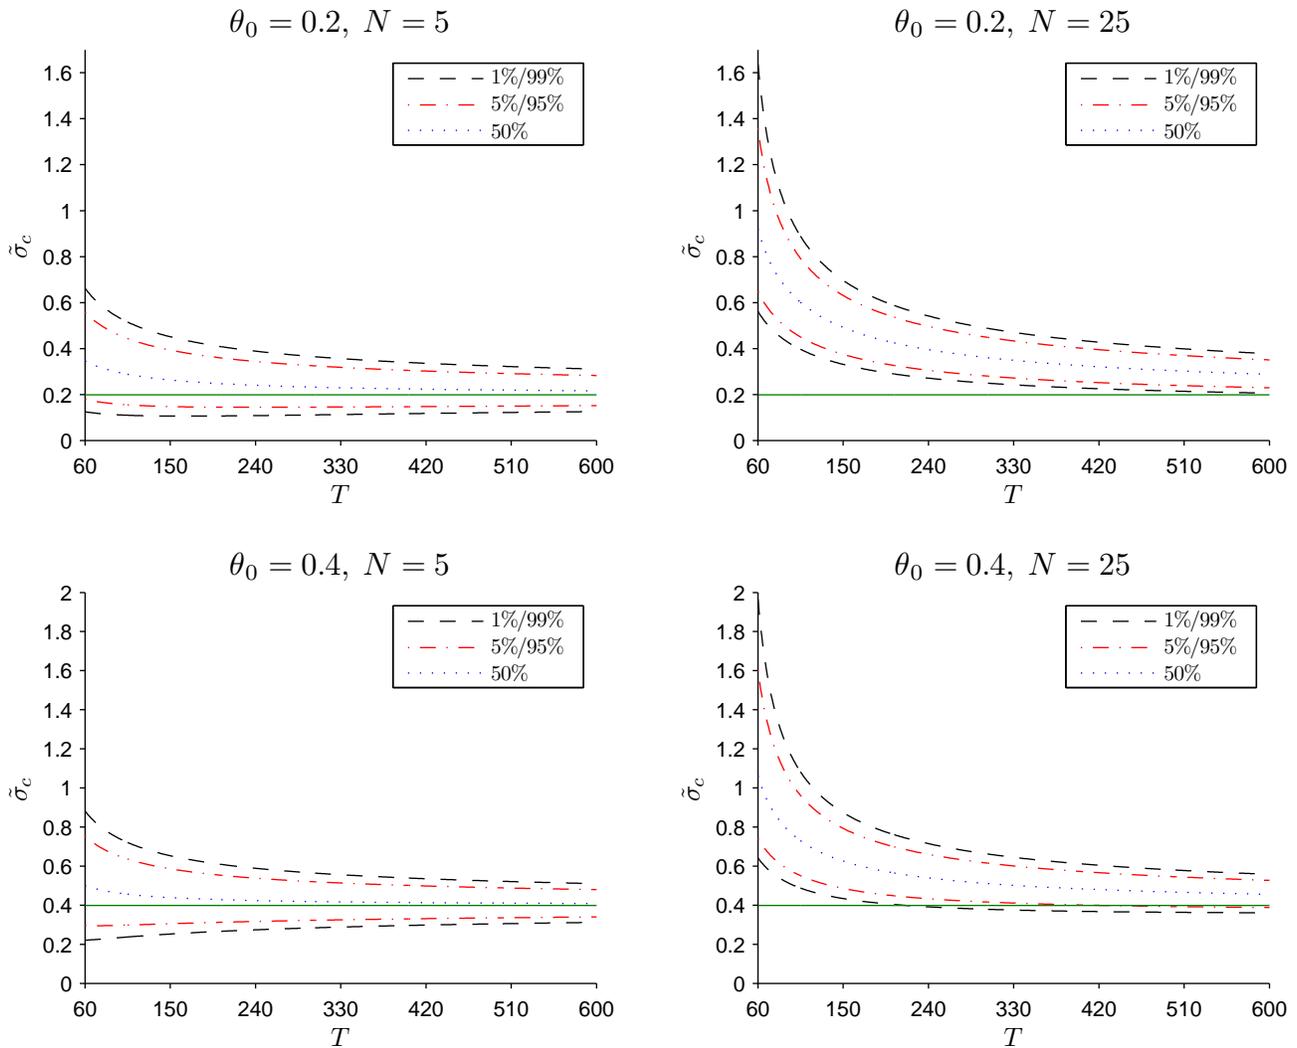


Figure 7

Exact distribution of the maximum likelihood estimator of the constrained Hansen-Jagannathan bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the maximum likelihood estimator of the constrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets. The gross risk-free rate is assumed to be 1.005 and the solid line in the figure represents the population value of the constrained HJ-bound.

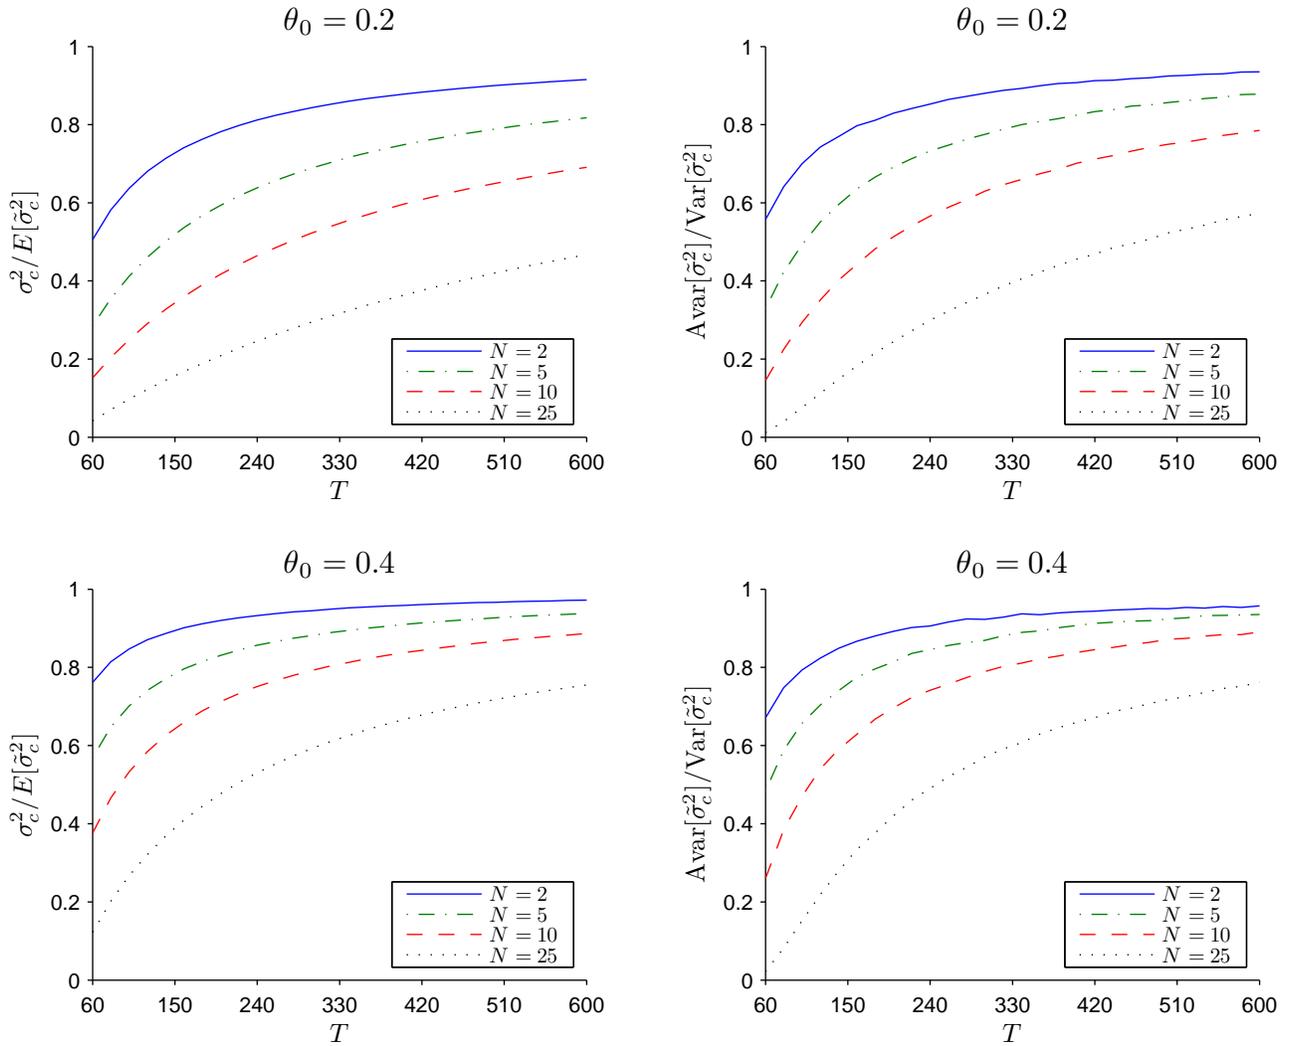


Figure 8

Ratio of asymptotic to exact moments of the maximum likelihood estimator of the constrained Hansen-Jagannathan bound. The figure presents the ratios of the asymptotic mean and variance to the exact mean and variance of the maximum likelihood estimator of the constrained HJ-bound ($\tilde{\sigma}_c^2$) for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The exact moments are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

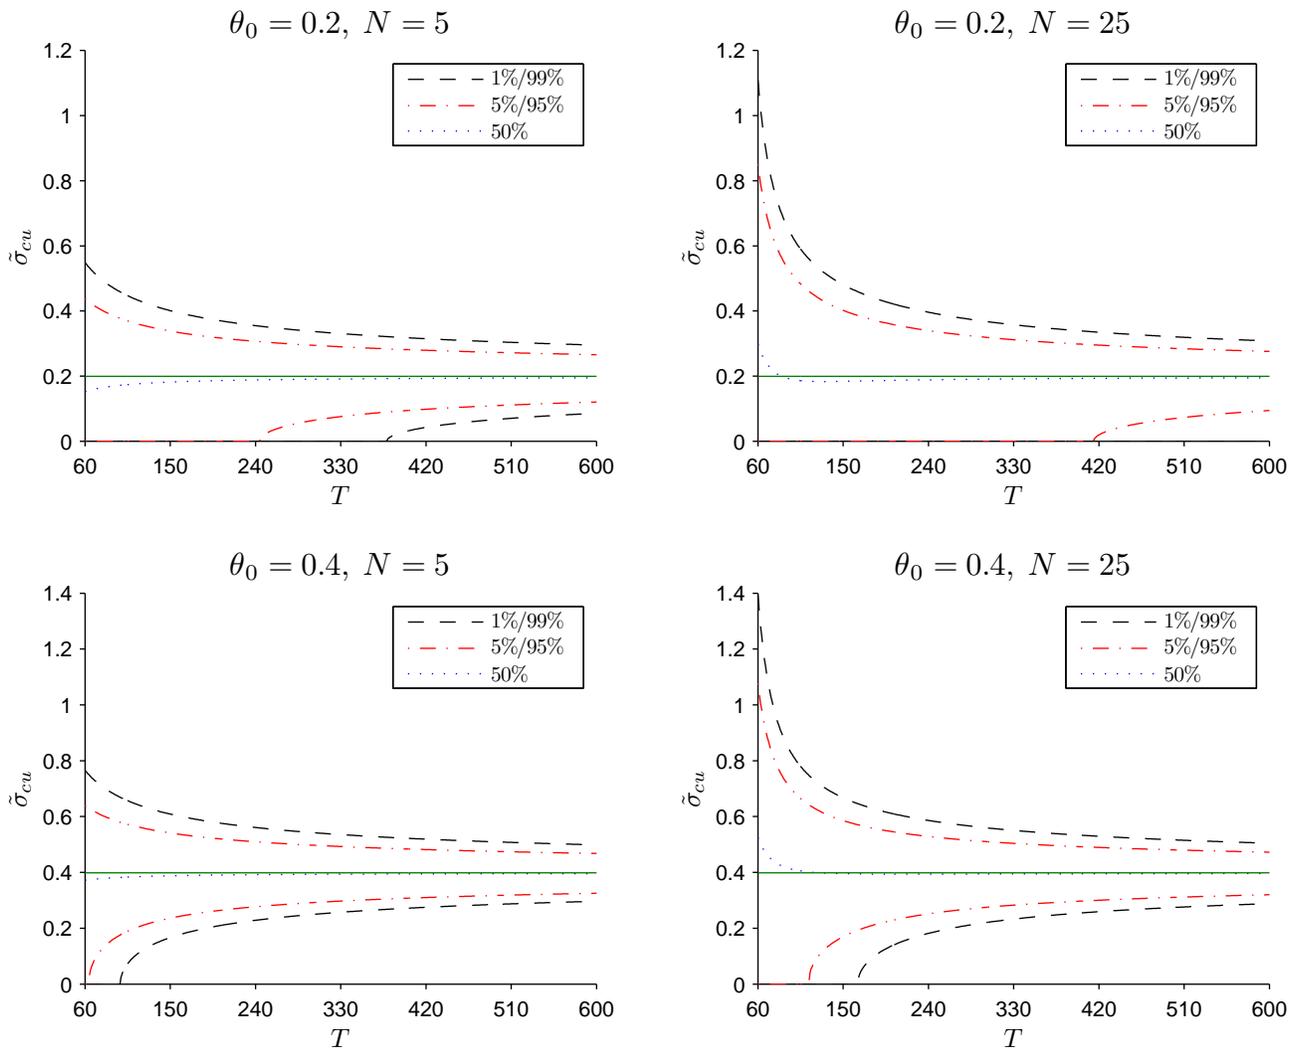


Figure 9

Finite sample distribution of an approximate unbiased estimator of the constrained Hansen-Jagannathan bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the finite sample distribution of an approximate unbiased estimator of the constrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets. The gross risk-free rate is assumed to be 1.005 and the solid line in the figure represents the population value of the constrained HJ-bound.

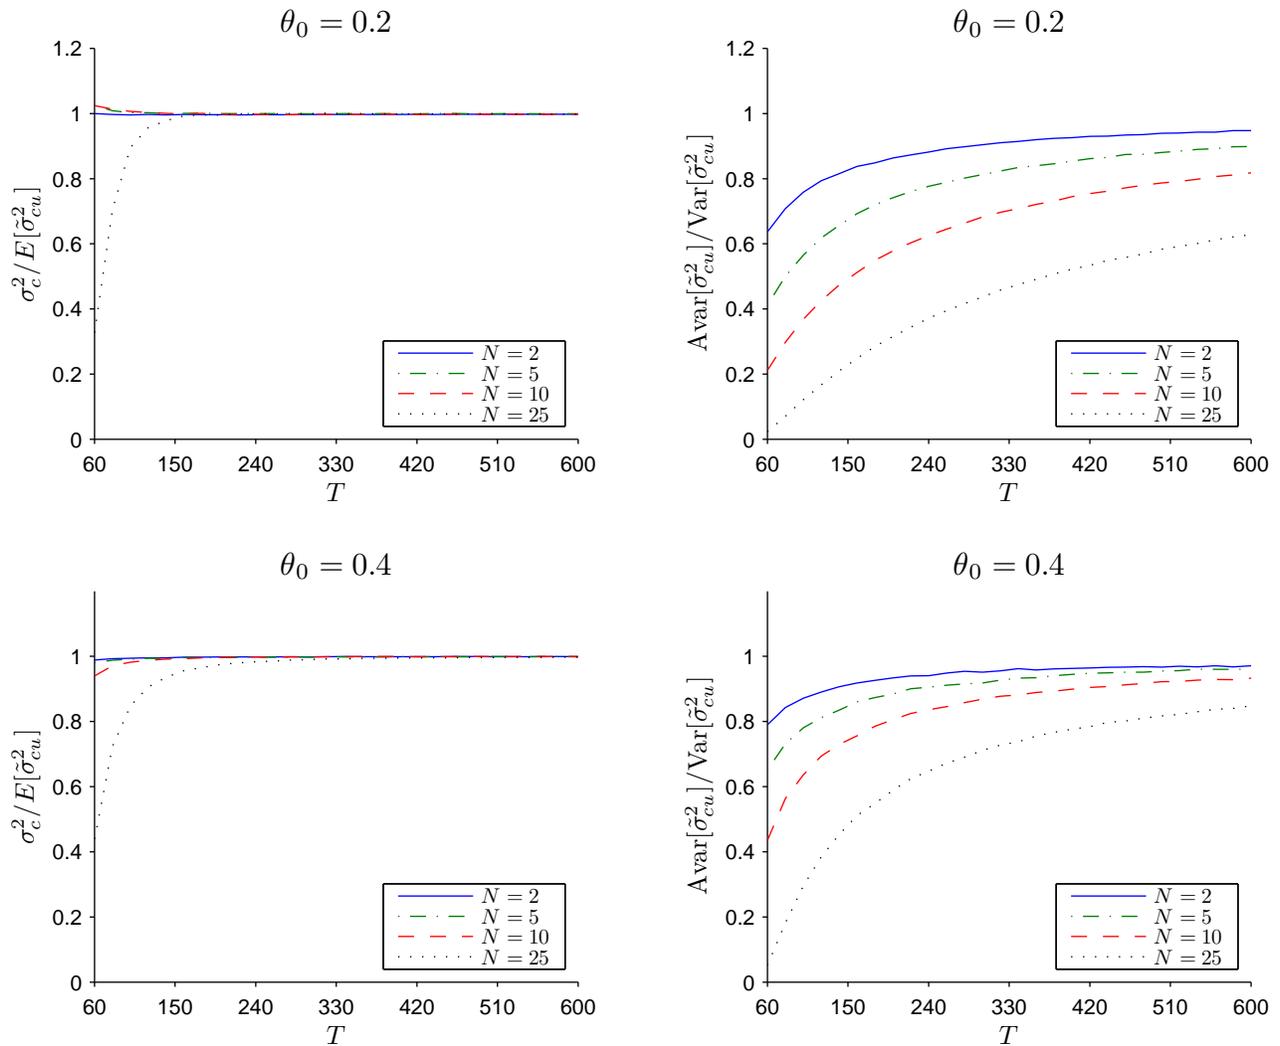


Figure 10

Ratio of asymptotic to finite moments of an approximate unbiased estimator of the constrained Hansen-Jagannathan bound. The figure presents the ratios of the asymptotic mean and variance to the finite mean and variance of an approximate unbiased estimator of the constrained HJ-bound ($\tilde{\sigma}_{cu}^2$) for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The finite moments are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

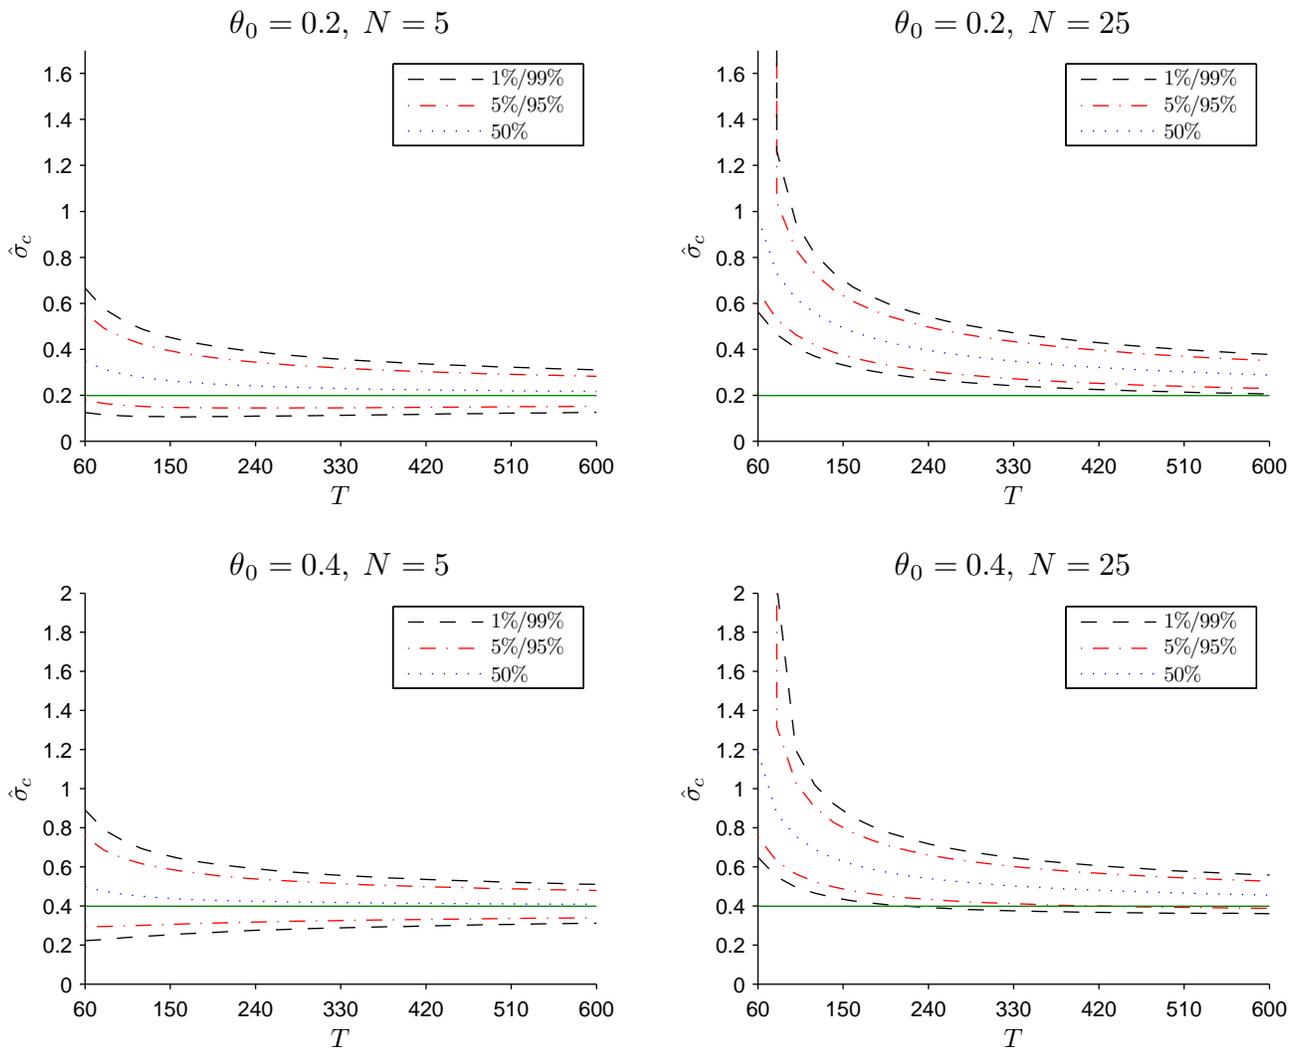


Figure 11

Exact distribution of the nonparametric estimator of the constrained Hansen-Jagannathan bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the sample estimator of the constrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the normality assumption. The exact moments are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets. The gross risk-free rate is assumed to be 1.005 and the solid line in the figure represents the population value of the constrained HJ-bound.

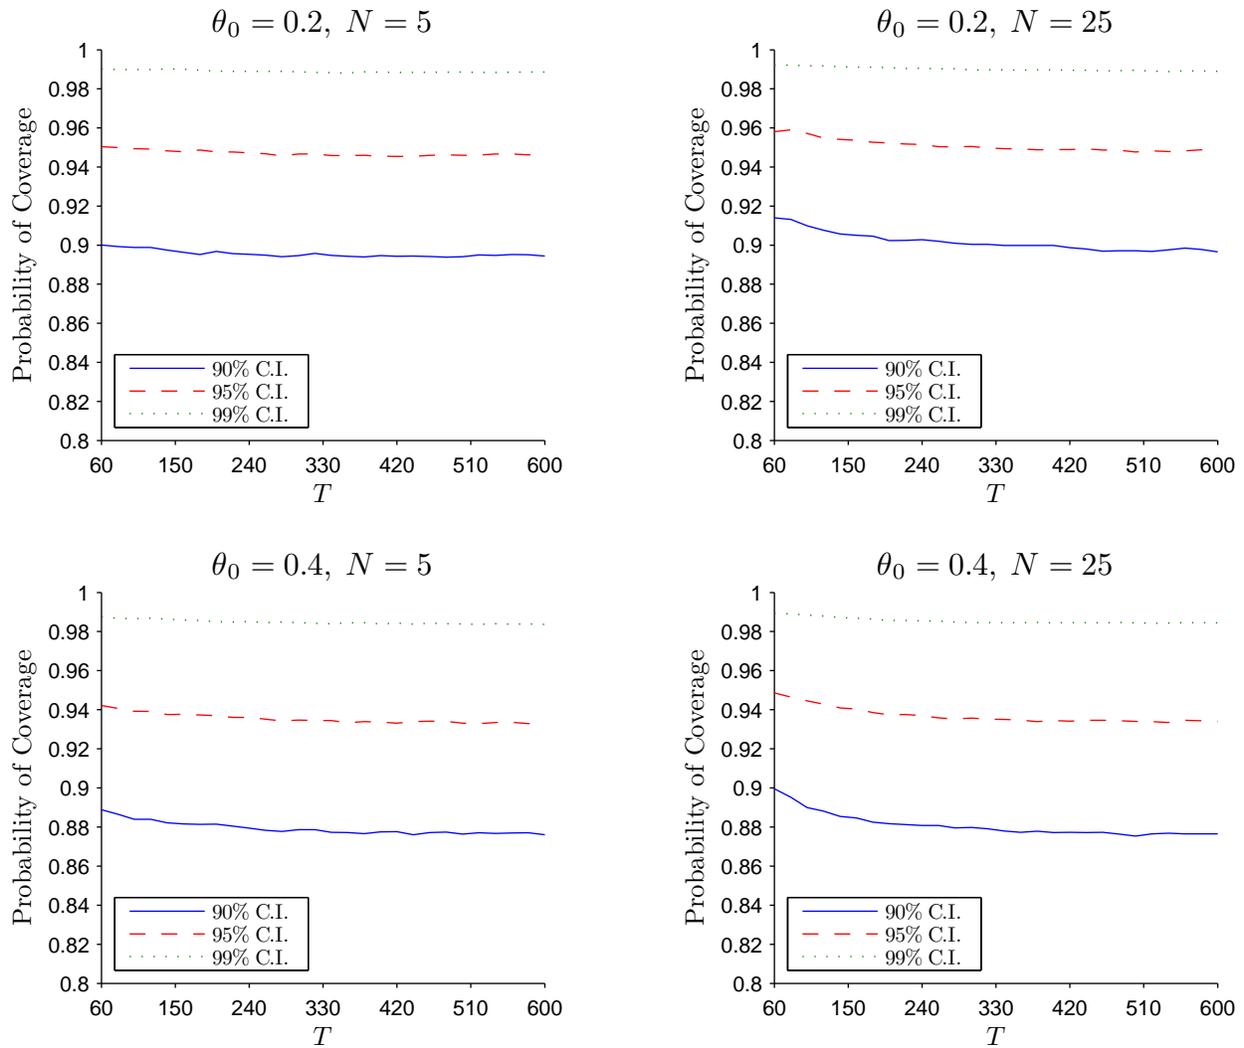


Figure 12

Probabilities of coverage of the confidence intervals for the unconstrained Hansen-Jagannathan bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the unconstrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the assumption that returns are multivariate t distributed with five degrees of freedom. The probabilities of coverage are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

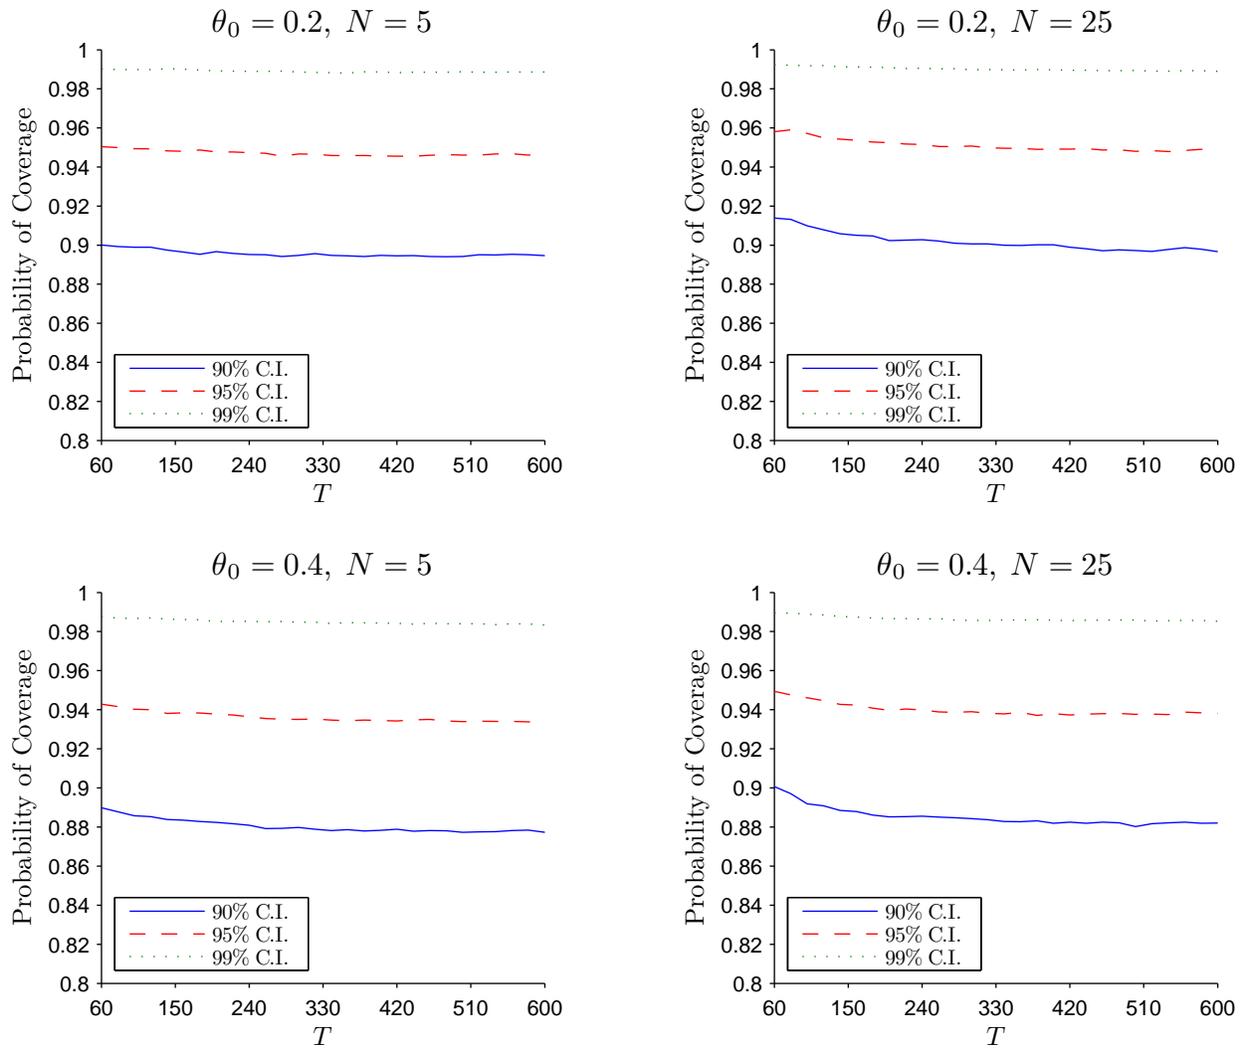


Figure 13

Probabilities of coverage of the confidence intervals for the constrained Hansen-Jagannathan bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the constrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the assumption that returns are multivariate t distributed with five degrees of freedom. The probabilities of coverage are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

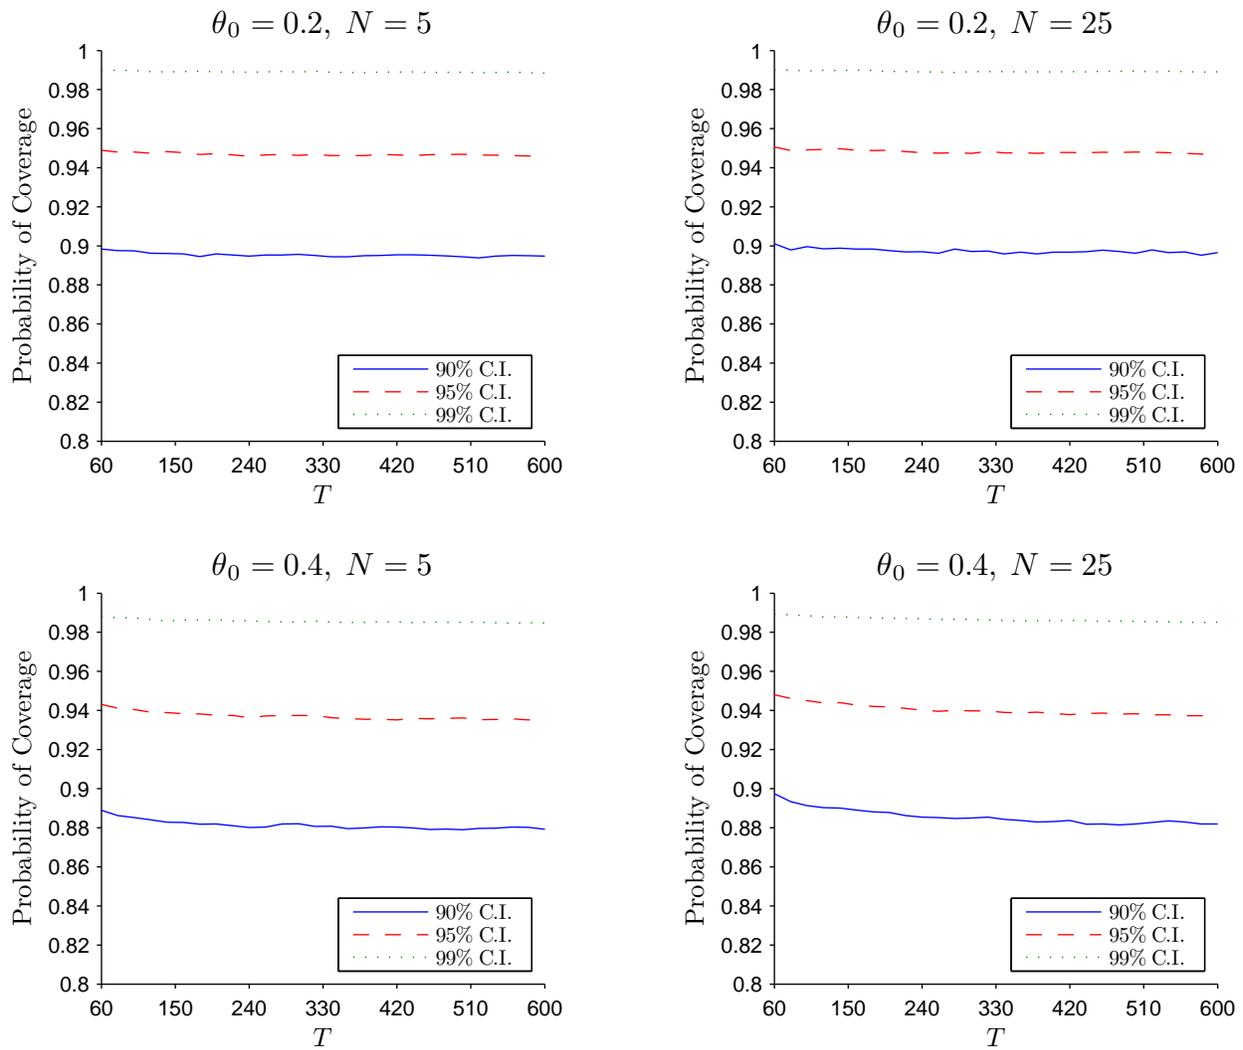


Figure 14

Probabilities of coverage of the confidence intervals for the unconstrained Hansen-Jagannathan bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the unconstrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the assumption that the transformed excess returns are GARCH(1,1) distributed, with the parameters chosen based on the monthly excess returns on the 25 Fama-French size and book-to-market ranked portfolios over the period 1946/1–2006/12. The probabilities of coverage are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.

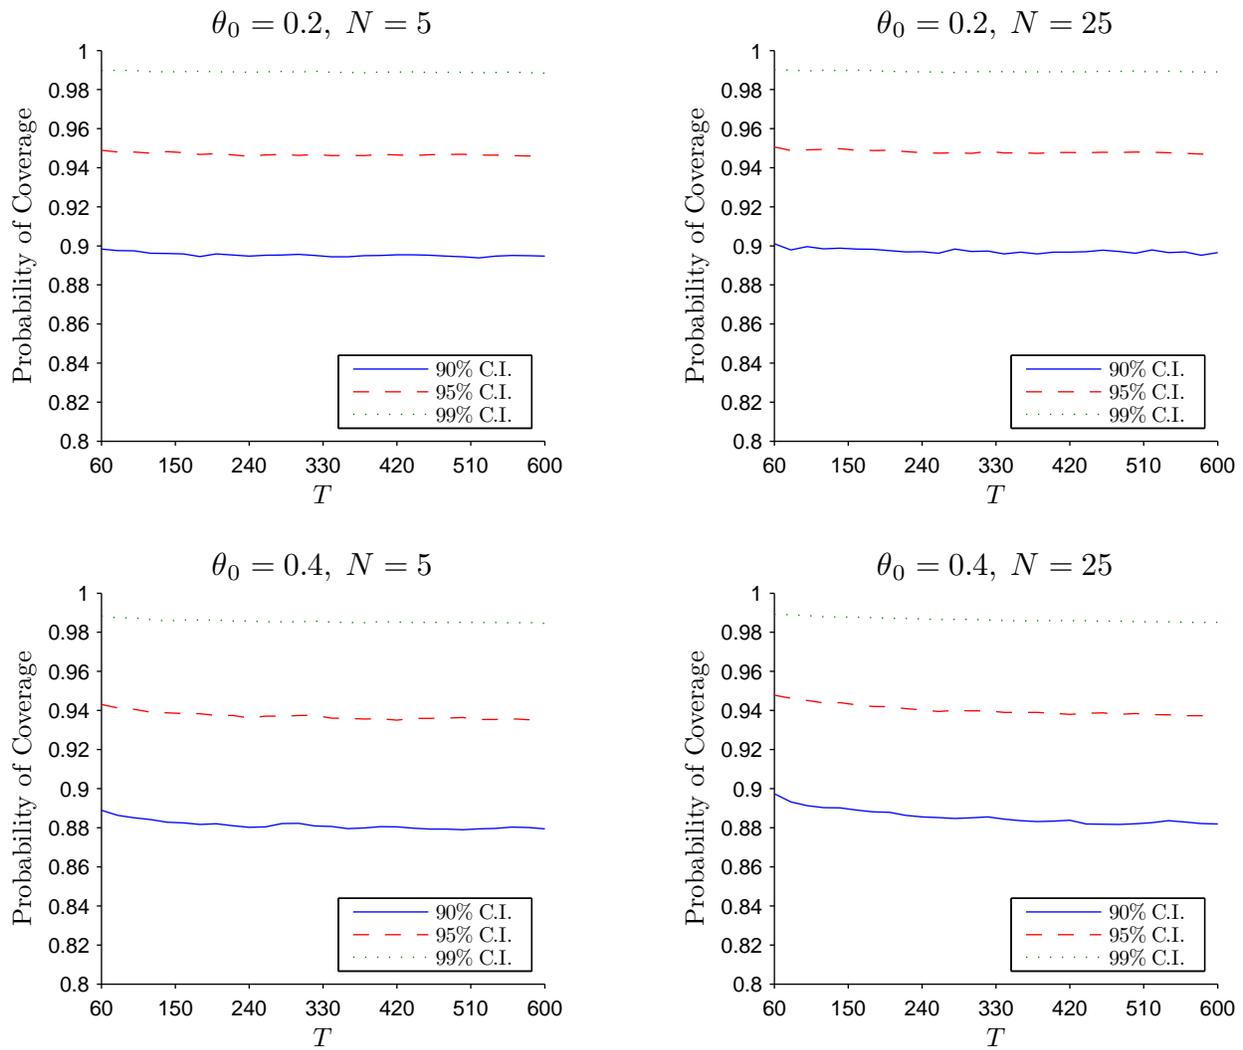


Figure 15

Probabilities of coverage of the confidence intervals for the constrained Hansen-Jagannathan bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the constrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under the assumption that the transformed excess returns are GARCH(1,1) distributed, with the parameters chosen based on the monthly excess returns on the 25 Fama-French size and book-to-market ranked portfolios over the period 1946/1–2006/12. The probabilities of coverage are estimated based on 100,000 simulations. The upper two panels are for $\theta_0 = 0.2$ and the lower two panels are for $\theta_0 = 0.4$, where θ_0 is the Sharpe ratio of the tangency portfolio of the N risky assets.