

The Exact Distribution of the Hansen-Jagannathan Bound

Raymond Kan and Cesare Robotti*

March 2015

JEL # C12, C13, G12

*We thank Jerome Detemple (the Editor), Wayne Ferson, Burton Hollifield, Sergei Sarkissian, Tim Simin, Andrea Tamoni, Jianfeng Yu, Chu Zhang, Guofu Zhou, an anonymous Associate Editor, three anonymous referees, and seminar participants at the Federal Reserve Bank of Atlanta, the 2008 China International Conference in Finance, and the 2008 Northern Finance Meetings for helpful discussions and comments. Kan gratefully acknowledges financial support from the National Bank Financial of Canada. Send correspondence to Cesare Robotti, Imperial College Business School, Imperial College London, Tanaka Building, South Kensington Campus, London SW7 2AZ, United Kingdom; telephone: +44 20 7594 2682. E-mail: c.robotti@imperial.ac.uk.

The Exact Distribution of the Hansen-Jagannathan Bound

Abstract

Under the assumption of multivariate normality of asset returns, this paper presents a geometric interpretation and the finite-sample distributions of the sample Hansen-Jagannathan 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.

1. Introduction

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.¹ Although the constrained HJ-bound is sharper than the unconstrained HJ-bound and is theoretically appealing, it 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 Snow (1991), Cecchetti, Lam, and Mark (1994), Burnside (1994), He and Modest (1995), and Hagiwara and Herce (1997).

Figure 1 plots the sample unconstrained HJ-bound and the nonparametric estimator of the constrained HJ-bound used in the existing literature. The sample HJ-bounds are computed using annual real non-overlapping returns on four size and book-to-market ranked portfolios, five industry portfolios, and a three-month Treasury bill over the post-war period (1952–2012).² The figure shows that for this choice of portfolio returns and sample period, there is a nontrivial difference between the two HJ-bounds. The constrained HJ-bound is markedly tighter than its unconstrained counterpart and potentially represents a higher hurdle for a proposed SDF.

Figure 1 about here

Since the two HJ-bounds can exhibit large differences, our objective in this paper is to understand what drives a wedge between the unconstrained and constrained HJ-bounds. We tackle

¹Examples include Bekaert and Hodrick (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.

²This set of portfolio returns has been shown to pose great challenges to existing asset-pricing models. The value-weighted equity returns are from Kenneth French's website. The five industries are (1) consumer durables, nondurables, wholesale, retail, and some services (laundries and repair shops), (2) manufacturing, energy, and utilities, (3) business equipment, telephone and television transmission, (4) healthcare, medical equipment, and drugs, and (5) mines, construction, building materials, transportation, hotels, business services, entertainment, and finance. The annual return on the three-month Treasury bill is obtained by compounding the one-month holding-period returns on a three-month bill, and the data is from the CRSP US Treasuries and Inflation Indices File. Real returns are obtained by deflating nominal returns by the Consumer Price Index (all urban consumers, not seasonally adjusted) from the Bureau of Labor Statistics.

this question by first examining the population properties of the two HJ-bounds and then characterizing their sampling distributions. The first part of the paper is concerned with the geometry of the unconstrained and constrained HJ-bounds. While there is a well-known mapping between the unconstrained HJ-bound and the mean-variance frontier of the portfolio returns, the mapping between the constrained HJ-bound and the mean-variance frontier has not been explored in depth. We contribute to an understanding of this relation by establishing an explicit link between the constrained HJ-bound and the Sharpe ratio of the portfolio with minimum second moment of truncated returns. Importantly, under a multivariate normality assumption on the portfolio returns, we are able to plot the mean-variance frontier of the truncated returns and analyze the differences between its shape and the shape of the mean-variance frontier. Formulating the constrained HJ-bound problem as a portfolio problem provides important insights. First, our analysis suggests that the squared Sharpe ratio of the tangency portfolio in the mean-variance frontier of the truncated portfolio returns is the primary driver of the constrained HJ-bound, and shows that a candidate SDF has to satisfy more stringent bounds when the squared Sharpe ratio of the tangency portfolio is high. This is consistent with the original goal of Hansen and Jagannathan (1991) of determining which characteristics of the asset returns represent the biggest challenge to a proposed SDF. Second, the dual formulation of the constrained HJ-bound problem shows that meaningful differences between the constrained and unconstrained HJ-bounds only arise when the Sharpe ratio of the tangency portfolio is large. Finally, the dual formulation of the constrained HJ-bound problem allows us to show that the traditional sample constrained HJ-bound can take on the value of infinity with positive probability.

As we mentioned earlier, 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. When there are N assets, one has to solve N nonlinear equations to obtain the constrained HJ-bound. In this paper, we show that under the assumption that 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 sample unconstrained and constrained HJ-bounds under the multivariate normality assumption.³ 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 empirical example may help illustrate the importance of reporting confidence intervals for the HJ-bounds instead of only presenting their point estimates. Using the same return data as in Figure 1, in Figure 2 we provide a comparison of the SDFs implied by three state-of-the-art consumption-based asset-pricing models with the HJ-bounds.⁴ The three asset-pricing models are the long-run risks model of Bansal, Kiku, and Yaron (2012a, BKY henceforth), the external habit model of Campbell and Cochrane (1999, CC henceforth), and the rare disasters model of Nakamura, Steinsson, Barro, and Ursúa (2013, NSBU henceforth). These models have been shown to match closely the historical average annual real returns on the risk-free bond and the equity market. The mean-standard deviation pairs of the SDFs implied by BKY, CC, and NSBU are computed via simulation using the estimated parameters reported in Bansal, Kiku, and Yaron (2012b), Aldrich and Gallant (2011), and Nakamura, Steinsson, Barro, and Ursúa (2013), respectively.⁵ Since the model-implied SDFs are positive, it seems reasonable to compare them with the more demanding HJ-bound that imposes the nonnegativity constraint. We present two different estimators of the constrained HJ-bound. 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

³Gordon, Samson, and Carmichael (1995) characterize the finite-sample uncertainty associated with the sample unconstrained HJ-bound using Bayesian inference methods.

⁴Since we consider Figure 2 to be primarily diagnostic, we do not present confidence regions for the mean-standard deviation pairs of the SDFs or formally test whether the point estimates of the SDF lie outside the HJ-bounds. Cochrane and Hansen (1992), Burnside (1994), and Cecchetti, Lam, and Mark (1994) develop classical hypothesis tests based on the distance between a given SDF and the HJ-bounds. These studies find that the point estimates of the SDF plot outside the sample unconstrained and constrained HJ-bounds 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).

⁵The mean-standard deviation pairs of the various SDFs are taken from Table 3 of Favero, Ortu, Tamoni, and Yang (2014). We refer the readers to their paper for a thorough description of the models and the computational details.

this paper. For this particular example, the two estimators of the constrained HJ-bound are quite far apart, consistent with the fact that the sample constrained HJ-bound, similarly to its unconstrained counterpart, can have a substantial finite-sample upward bias. While the use of the maximum likelihood estimator of the constrained HJ-bound would lead us to conclude that all models are rejected by the data, one would reach a more optimistic conclusion for NSBU when using the approximate unbiased estimator of the constrained HJ-bound. In addition, comparing the SDF implied by NSBU with the HJ-bounds in Figure 1, it is immediately evident that this model satisfies the sample unconstrained HJ-bound but does not meet the tighter sample constrained HJ-bound. Therefore, at least for this empirical example, the magnitude of the bias appears to be economically significant, and the use of the two HJ-bounds leads to different conclusions on the performance of a given SDF. Since the approximate unbiased estimator can still be quite volatile, to obtain a better idea of where the population HJ-bound may actually lie, we construct 95% confidence intervals for the constrained HJ-bound (dotted lines) using the methodology described later in the paper. The confidence intervals in Figure 2 are quite wide, indicating that there is substantial uncertainty about the exact location of the constrained HJ-bound.⁶ Taking sampling error into account, both NSBU and CC might not be entirely at odds with the data.⁷ This stands in sharp contrast to the rejections of the models that we obtain when merely relying on the sample constrained HJ-bound.

Figure 2 about here

The rest of the paper is organized as follows. Section 2 presents the unconstrained and constrained HJ-bounds and our main results on the constrained HJ-bound under the assumption that the returns are multivariate normally distributed. In Section 3, 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 4, under the multivariate normality assumption, we present the finite-sample distributions of the sample unconstrained and constrained

⁶Gregory and Smith (1992) and Burnside (1994) show that the large sampling error in the sample HJ-bounds can lead to a false rejection of the model. In related work, Britten-Jones (1999) shows the importance and magnitude of sampling error in estimates of the weights of a mean-variance efficient portfolio. The main message that emerges from these studies is that merely relying on point estimates of HJ-bounds or tangency portfolio weights can be misleading.

⁷Instead of constructing confidence intervals for the HJ-bounds, Burnside (1994) constructs confidence regions for the parameters of the candidate SDF. His 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.

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 lemmas and propositions.

2. Population Hansen-Jagannathan Bounds

In this section, we summarize existing results and present some new ones on the HJ-bounds. The 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 the 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.⁸ 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.

⁸Although 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).

2.1. Unconstrained Hansen-Jagannathan Bound

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

$$E[mR] = 1_N, \quad (1)$$

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

$$R_p = (1 - w'1_N)R_0 + w'R = R_0 + w'(R - R_01_N). \quad (2)$$

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

$$\min_w E[(R_0 + w'(R - R_01_N))^2]. \quad (3)$$

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

$$w^* = -\frac{R_0}{1 + \theta_0^2} V^{-1}(\mu - R_01_N), \quad (4)$$

where

$$\theta_0^2 = (\mu - R_01_N)'V^{-1}(\mu - R_01_N) = a - 2bR_0 + cR_0^2, \quad (5)$$

and $a = \mu'V^{-1}\mu$, $b = 1'_N V^{-1}\mu$, and $c = 1'_N V^{-1}1_N$ 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 $V^{-1}(\mu - R_01_N)/(b - cR_0)$ 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

$$R_{p^*} = R_0 + w^{*'}(R - R_01_N). \quad (6)$$

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

$$m_0 = \frac{R_{p^*}}{\|R_{p^*}\|^2} = \frac{1 - (\mu - R_0\mathbf{1}_N)'V^{-1}(R - \mu)}{R_0}, \quad (7)$$

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 (7), we have (1) $E[m_0] = 1/R_0$, (2) $E[m_0R] = \mathbf{1}_N$, (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

$$\text{Var}[m_0] = \frac{(\mu - R_0\mathbf{1}_N)'V^{-1}(\mu - R_0\mathbf{1}_N)}{R_0^2} = a\mu_m^2 - 2b\mu_m + c \equiv \sigma_0^2, \quad (8)$$

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 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 given by

$$\theta_0 = \frac{\sigma_{p^*}}{\mu_{p^*}} = \frac{\sigma_0}{\mu_m}. \quad (9)$$

In Figure 3, we provide a graphical illustration of this relation in the (σ_p, μ_p) space. 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. 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

2.2. 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. To obtain a nonnegative minimum-variance SDF, Hansen and Jagannathan (1991) propose to first solve the dual portfolio problem

$$\min_w E[\max[0, R_0 + w'(R - R_0\mathbf{1}_N)]^2]. \quad (10)$$

The problem amounts to finding a portfolio that minimizes the second moment of its truncated return, i.e., $\min_w E[R_p^+{}^2]$, where $R_p^+ = \max[0, R_p]$. This portfolio problem is nontrivial to solve because the second moment of a truncated portfolio return 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 (as it is the case, for example, under multivariate elliptical distributions), knowing the mean and the covariance matrix of R is in general not sufficient for us to solve this problem.

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

$$m_c = \frac{R_{q^*}^+}{\|R_{q^*}^+\|^2}. \quad (11)$$

Lemma 2 summarizes the properties of m_c .

Lemma 2. *For m_c defined in (11), we have (1) $E[m_c] = 1/R_0$, (2) $E[m_c R] = \mathbf{1}_N$, (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 $\text{Cov}[m, m_c] \geq \text{Var}[m_c] = \sigma_c^2$ and $\text{Var}[m] \geq \sigma_c^2$.*

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 σ_c^2 provides a lower bound for the variance of all admissible SDFs that are nonnegative.⁹ In many ways, Lemma 2 is almost identical to Lemma 1 after we replace m_0 with m_c and R_{p^*} with $R_{q^*}^+$. 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 the inequality $\text{Cov}[m, m_c] \geq \sigma_c^2$. The reason is that $R_{q^*}^+$ is not a portfolio return. As a result, we do not have the exact pricing result $E[mR_{q^*}^+] = 1$ but just the inequality $E[mR_{q^*}^+] \geq 1$.

There is a well known mapping between the unconstrained HJ-bound and the mean-variance frontier of the portfolio returns as given in (9). A similar mapping can also be developed for the constrained HJ-bound. Let $\mu_p^+ \equiv E[R_p^+]$ and $\sigma_p^+ \equiv (\text{Var}[R_p^+])^{\frac{1}{2}}$ be the mean and standard deviation of the truncated gross return on portfolio p , respectively. The following lemma presents the linkage between the mean-variance frontier of the truncated portfolio returns and the constrained HJ-bound.

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

$$\theta_c^2 = \frac{(E[R_{q^*}^+] - R_0)^2}{\text{Var}[R_{q^*}^+]}. \quad (12)$$

We have

$$\theta_c = \frac{\sigma_{q^*}^+}{\mu_{q^*}^+} = \frac{\sigma_c}{\mu_m}. \quad (13)$$

Comparing (13) with (9), 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 of the portfolio returns with the mean and variance of the truncated portfolio returns.

Although the results on the constrained and unconstrained HJ-bounds are quite similar, the constrained HJ-bound is rarely used in empirical work. We believe that part of the reason is that it is difficult to find the portfolio that minimizes the second moment of its truncated return in (10).

⁹Using Hölder's inequality and exploiting the fact that the SDF is assumed to be a positive random variable, Snow (1991) derives explicit restrictions on general noncentral moments of the SDF. Extending Hansen and Jagannathan (1991), Snow (1991) provides the greatest lower bound on the δ th moment of the SDF as a function of the price of some portfolio and the p th moment of the truncated payoff on that portfolio, where $1/\delta + 1/\rho = 1$.

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 second moment of its truncated return. As a result, we cannot plot the mean-variance frontier of the truncated portfolio returns, and it becomes hard to visualize the relation between this frontier and the constrained HJ-bound as described in Lemma 3. 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 for the portfolio that minimizes the second moment of its truncated return.¹⁰ 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.

2.3. 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 first and second moments of the truncated portfolio return, R_p^+ , which is the first step towards finding the portfolio that minimizes the second moment of its truncated return.

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 first and second moments of R_p^+ are given by*

$$E[R_p^+] = \mu_p \Phi(\eta) + \sigma_p \phi(\eta), \quad (14)$$

$$E[R_p^{+2}] = (\mu_p^2 + \sigma_p^2) \Phi(\eta) + \mu_p \sigma_p \phi(\eta), \quad (15)$$

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.

¹⁰Cecchetti, 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 for the constrained HJ-bound. In addition to the multivariate normality case, we also obtain an analytical solution for the constrained HJ-bound under a multivariate elliptical distribution assumption. The results of this analysis are available from the authors upon request.

In addition to Lemma 4, we need the following lemma to find the portfolio that minimizes the second moment of its truncated return.

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 $c > 0$.*

Using the results in Lemma 4 and Lemma 5, we present the explicit solution to the constrained HJ-bound problem in Proposition 1.

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

$$g(\eta) = \frac{1}{\theta_0}, \quad (16)$$

where θ_0 is defined in (5). Then, the portfolio that minimizes the second moment of its truncated return has the following weights in the N risky assets:

$$w^* = -\frac{R_0}{\theta_0(\eta^* + \theta_0)} V^{-1} (\mu - R_0 \mathbf{1}_N). \quad (17)$$

In addition, the constrained HJ-bound is given by

$$\text{Var}[m_c] = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)} - \frac{1}{R_0^2} = \frac{\sigma_0(\eta^* \mu_m + \sigma_0)}{\Phi(\eta^*)} - \mu_m^2 \equiv \sigma_c^2, \quad (18)$$

where $\mu_m = 1/R_0$ and σ_0^2 is the unconstrained HJ-bound defined in (8).

By examining (17), we can see that under the normality assumption, the portfolio that minimizes the second moment of its truncated return, just like the minimum second moment portfolio in (4), is a linear combination of the tangency portfolio and the risk-free asset. From (16), we know

$$\frac{1}{\theta_0} = \eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} > \eta^*, \quad (19)$$

and we have

$$0 < \theta_0(\eta^* + \theta_0) < 1 + \theta_0^2. \quad (20)$$

Consequently, the portfolio with minimum second moment of truncated returns in (17) involves short-selling more of the tangency portfolio than the minimum second moment portfolio in (4). With the nonnegativity constraint, σ_c^2 is naturally greater than σ_0^2 . From (16), 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 monotonically 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*

$$h(\theta_0) = \sigma_c^2 - \sigma_0^2, \quad \tilde{h}(\theta_0) = \frac{\sigma_c^2}{\sigma_0^2} - 1. \quad (21)$$

We have (1) $h(\theta_0)$ and $\tilde{h}(\theta_0)$ are positive and strictly increasing functions 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 suggests that when the Sharpe ratio of the tangency portfolio is small, the constrained and unconstrained HJ-bounds will be very close to each other. It is only when the tangency portfolio has a large Sharpe ratio that we can expect some meaningful differences between the two HJ-bounds. Note that the Sharpe ratio of the tangency portfolio is an increasing function of the number of test assets and of the return horizon. Therefore, we should expect a bigger difference between the unconstrained and constrained HJ-bounds when more test assets are employed or when the return horizon is long.

Proposition 1 provides a useful characterization of the portfolio q^* with minimum second moment of truncated returns and allows us to better understand the relation between the Sharpe ratio of q^* and the constrained HJ-bound in Lemma 3. In Figure 4, we provide a graphical illustration of this relation in the (σ_p^+, μ_p^+) space under the assumption that the 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 variance frontier of the truncated portfolio returns. It can be readily shown that just like the mean-variance frontier of the returns, the mean-variance frontier of the truncated returns is 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 of the returns and the mean-variance frontier of the truncated returns. The first one is that the mean-variance frontier of the truncated returns is not represented by two straight lines as it is the case for the mean-variance

frontier of the returns. When the weight of the risk-free asset in a portfolio is close to one (i.e., near the point $(0, R_0)$), the gross return on the portfolio has a very small probability of assuming a negative value, and the mean-variance frontier of the truncated returns is almost identical to the mean-variance frontier of the returns. When the portfolio has significant positive or negative weights in the tangency portfolio, then the mean and standard deviation of the truncated portfolio return can differ significantly from the mean and standard deviation of the portfolio return. The second difference is that the mean of the truncated portfolio return has a lower bound, so the lower curve does not continue to go down and instead turns around after reaching a minimum.¹¹ Once it turns around, the curve becomes the maximum variance frontier of the truncated returns rather than the minimum variance frontier of the truncated returns. 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 expected truncated portfolio return, since the negative returns are dropped in the calculation of the mean of the truncated portfolio return. Since $E[(R_p^+)^2] = (\mu_p^+)^2 + (\sigma_p^+)^2$, the portfolio with minimum second moment of truncated returns has the shortest distance from the origin. To locate q^* , we draw a circle with its center at the origin, and the location of q^* can be obtained from the point where the circle is tangent to the minimum variance frontier of the truncated portfolio returns. Note that the absolute value of the Sharpe ratio of q^* as well as $\sigma_{q^*}^+/\mu_{q^*}^+$ 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

3. 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 , $t = 1, \dots, T$. The

¹¹Let u^* be the solution to the equation $\phi(u)/\Phi(u) = \theta_0$. It can be shown that $\min_w \mu_p^+ = R_0\Phi(u^*)$.

mean and covariance matrix of R_t can be estimated using the sample quantities

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T R_t, \quad (22)$$

$$\hat{V} = \frac{1}{T} \sum_{t=1}^T (R_t - \hat{\mu})(R_t - \hat{\mu})'. \quad (23)$$

To estimate the unconstrained HJ-bound for a given value of $\mu_m = 1/R_0$, researchers typically use the sample counterpart of (8)

$$\hat{\sigma}_0^2 = \hat{a}\mu_m^2 - 2\hat{b}\mu_m + \hat{c}, \quad (24)$$

where \hat{a} , \hat{b} , and \hat{c} are given by

$$\hat{a} = \hat{\mu}'\hat{V}^{-1}\hat{\mu}, \quad \hat{b} = 1'_N\hat{V}^{-1}\hat{\mu}, \quad \hat{c} = 1'_N\hat{V}^{-1}1_N. \quad (25)$$

For the constrained HJ-bound, Hansen, Heaton, and Luttmer (1995) suggest an estimator of σ_c^2 that can be obtained in two steps. We first estimate $E[R_{q^*}^{+2}]$. This can be accomplished by computing the sample counterpart of (10)

$$\hat{\lambda} = \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 1_N)]^2, \quad (26)$$

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

$$\hat{\sigma}_c^2 = \frac{1}{\hat{\lambda}} - \mu_m^2. \quad (27)$$

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 (26) 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 (17) 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

$$g(\eta) = \frac{\mu_m}{\hat{\sigma}_0}, \quad (28)$$

where $\hat{\sigma}_0^2$ is the sample unconstrained HJ-bound defined in (24). Using $\hat{\eta}^*$, we compute the maximum likelihood estimator of σ_c^2 as

$$\tilde{\sigma}_c^2 = \frac{\hat{\sigma}_0(\hat{\eta}^* \mu_m + \hat{\sigma}_0)}{\Phi(\hat{\eta}^*)} - \mu_m^2. \quad (29)$$

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 (28), so 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 this estimator can be very volatile. 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 noisy. The finite-sample performance of these two estimators under normal and nonnormal distributions will be studied in Section 4.

3.1. 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 variances 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

$$m_{0,t} = \frac{1 - (\mu - R_0 \mathbf{1}_N)' V^{-1} (R_t - \mu)}{R_0} = \mu_m - (\mu_m \mu - \mathbf{1}_N)' V^{-1} (R_t - \mu), \quad (30)$$

$$m_{c,t} = \frac{R_{q^*,t}^+}{\|R_{q^*,t}^+\|^2}. \quad (31)$$

With some simplifications, the asymptotic results of Hansen, Heaton, and Luttmer (1995) can be written as

$$\sqrt{T}(\hat{\sigma}_0^2 - \sigma_0^2) \overset{A}{\sim} N(0, v_0), \quad (32)$$

$$\sqrt{T}(\hat{\sigma}_c^2 - \sigma_c^2) \overset{A}{\sim} N(0, v_c), \quad (33)$$

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

$$\phi_{0,t} = m_{0,t}^2 - \mu_m^2 - \sigma_0^2, \quad (34)$$

$$\phi_{c,t} = m_{c,t}^2 - \mu_m^2 - \sigma_c^2. \quad (35)$$

To obtain a consistent estimator of v_0 , we can replace $\phi_{0,t}$ with

$$\hat{\phi}_{0,t} = \hat{m}_{0,t}^2 - \mu_m^2 - \hat{\sigma}_0^2, \quad (36)$$

where

$$\hat{m}_{0,t} = \frac{1 - (\hat{\mu} - R_0 \mathbf{1}_N)' \hat{V}^{-1} (R_t - \hat{\mu})}{R_0} = \mu_m - (\mu_m \hat{\mu} - \mathbf{1}_N)' \hat{V}^{-1} (R_t - \hat{\mu}) \quad (37)$$

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

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

where

$$\hat{m}_{c,t} = \frac{R_{\hat{q}^*,t}^+}{\frac{1}{T} \sum_{t=1}^T (R_{\hat{q}^*,t}^+)^2} = \frac{R_{\hat{q}^*,t}^+}{R_0 \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+}, \quad (39)$$

and $R_{\hat{q}^*,t}^+$ is the gross return on the portfolio with sample minimum second moment of truncated returns at time t .¹² When returns are i.i.d. multivariate normally distributed, we can derive analytical expressions of 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.

¹²The last equality in (39) follows from the fact that the first order condition of the minimization problem is $\frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ (R_t - R_0 \mathbf{1}_N) = 0_N$, which implies that

$$\frac{1}{T} \sum_{t=1}^T (R_{\hat{q}^*,t}^+)^2 = \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ R_{\hat{q}^*,t} = \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+ [R_0 + w^{*'} (R_t - R_0 \mathbf{1}_N)] = R_0 \frac{1}{T} \sum_{t=1}^T R_{\hat{q}^*,t}^+. \quad (40)$$

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*

$$\sqrt{T}(\hat{\sigma}_0^2 - \sigma_0^2) \overset{A}{\approx} N\left(0, \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4}\right), \quad (41)$$

$$\sqrt{T}(\hat{\sigma}_c^2 - \sigma_c^2) \overset{A}{\approx} N\left(0, \frac{\theta_0^3[\theta_0(\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5)]}{R_0^4\Phi(\eta^*)^3} - \frac{\theta_0^2(\eta^* + \theta_0)^2}{R_0^4\Phi(\eta^*)^2}\right), \quad (42)$$

$$\sqrt{T}(\tilde{\sigma}_c^2 - \sigma_c^2) \overset{A}{\approx} N\left(0, \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4\Phi(\eta^*)^2}\right), \quad (43)$$

where θ_0^2 is defined in (5) 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 $\tilde{\sigma}_c^2$.

4. 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 $\hat{\sigma}_c^2$ 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 sample squared Sharpe ratio of the tangency portfolio

$$\hat{\theta}_0^2 = (\hat{\mu} - R_0\mathbf{1}_N)' \hat{V}^{-1} (\hat{\mu} - R_0\mathbf{1}_N) \quad (44)$$

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

$$\hat{\theta}_0^2 \sim \left(\frac{N}{T-N} \right) F_{N, T-N}(T\theta_0^2), \quad (45)$$

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

$$E[\hat{\theta}_0^r] = \frac{\Gamma\left(\frac{N+r}{2}\right) \Gamma\left(\frac{T-N-r}{2}\right)}{\Gamma\left(\frac{N}{2}\right) \Gamma\left(\frac{T-N}{2}\right)} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right), \quad (46)$$

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

$$E[\hat{\theta}_0^r] = \frac{\left(\frac{N}{2}\right)_{r/2}}{\left(\frac{T-N-r}{2}\right)_{r/2}} \sum_{i=0}^{r/2} \frac{\binom{r/2}{i} \left(\frac{T\theta_0^2}{2}\right)^i}{\left(\frac{N}{2}\right)_i}, \quad (47)$$

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

4.1. 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

$$P[\hat{\sigma}_0^2 < v] = P[\hat{\theta}_0^2 < R_0^2 v] = F_{N, T-N, T\theta_0^2}\left(\frac{(T-N)R_0^2 v}{N}\right), \quad (48)$$

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.¹³ We then plot the 1st, 5th, 50th, 95th, and 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

¹³Note 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%.

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 the 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 (47), we can obtain the exact mean and variance of the sample unconstrained HJ-bound $\hat{\sigma}_0^2$ as¹⁴

$$E[\hat{\sigma}_0^2] = \frac{N + T\theta_0^2}{(T - N - 2)R_0^2} \quad \text{if } T - N > 2, \quad (50)$$

$$\text{Var}[\hat{\sigma}_0^2] = \frac{2[(N + T\theta_0^2)^2 + (N + 2T\theta_0^2)(T - N - 2)]}{(T - N - 2)^2(T - N - 4)R_0^4} \quad \text{if } T - N > 4. \quad (51)$$

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

$$\frac{\sigma_0^2}{E[\hat{\sigma}_0^2]} = \frac{(T - N - 2)\theta_0^2}{N + T\theta_0^2}, \quad (52)$$

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

¹⁴Using 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):

$$\hat{\sigma}_{0u}^2 = \left(\frac{T - N - 2}{T} \right) \hat{\sigma}_0^2 - \frac{N}{T} \mu_m^2. \quad (49)$$

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

4.2. 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 (29) as

$$\tilde{\sigma}_c^2 = \frac{1}{R_0^2} \left[\frac{\hat{\theta}_0(\hat{\eta}^* + \hat{\theta}_0)}{\Phi(\hat{\eta}^*)} - 1 \right], \quad (54)$$

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 monotonically increasing function of $\hat{\theta}_0^2$. Denoting the monotonic relation between $\tilde{\sigma}_c^2$ and $\hat{\theta}_0^2$ by

$$\tilde{\sigma}_c^2 = f(\hat{\theta}_0^2), \quad \hat{\theta}_0^2 = f^{-1}(\tilde{\sigma}_c^2), \quad (55)$$

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

$$P[\tilde{\sigma}_c^2 < v] = P[f^{-1}(\tilde{\sigma}_c^2) < f^{-1}(v)] = F_{N,T-N,T\theta_0^2} \left(\frac{(T-N)f^{-1}(v)}{N} \right). \quad (56)$$

The only difference is that in this case we need to numerically compute $f^{-1}(v)$. Since $f(x)$ is a monotonically 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 $R_0 = 1.005$ 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

$$\tilde{\sigma}_c^2 = f(\hat{\theta}_0^2) \approx f(\theta_0^2) + f'(\theta_0^2)(\hat{\theta}_0^2 - \theta_0^2) = \sigma_c^2 + \frac{1}{R_0^2\Phi(\eta^*)}(\hat{\theta}_0^2 - \theta_0^2). \quad (57)$$

Taking the expectation of both sides, we obtain

$$E[\tilde{\sigma}_c^2] \approx \sigma_c^2 + \frac{1}{R_0^2 \Phi(\eta^*)} (E[\hat{\theta}_0^2] - \theta_0^2) = \sigma_c^2 + \frac{N + (N + 2)\theta_0^2}{(T - N - 2)R_0^2 \Phi(\eta^*)}. \quad (58)$$

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

$$\tilde{\sigma}_c^2 - \frac{N + (N + 2)\hat{\theta}_0^2}{(T - N - 2)R_0^2 \Phi(\hat{\eta}^*)} \quad (59)$$

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

$$\tilde{\sigma}_{cu}^2 = \max \left[0, \tilde{\sigma}_c^2 - \frac{N + (N + 2)\hat{\theta}_{0u}^2}{(T - N - 2)R_0^2 \Phi(\hat{\eta}_u^*)} \right] \quad (60)$$

as an adjusted estimator of σ_c^2 , where

$$\hat{\theta}_{0u}^2 = \max \left[0, \left(\frac{T - N - 2}{T} \right) \hat{\theta}_0^2 - \frac{N}{T} \right] \quad (61)$$

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

$$u + \frac{\phi(u)}{\Phi(u)} = \frac{1}{\hat{\theta}_{0u}}. \quad (62)$$

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 ($N = 2, 5, 10, \text{ and } 25$) 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 $\tilde{\sigma}_c^2$, 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

4.3. 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 $\hat{\sigma}_c^2$ coincides with the distribution of*

$$\frac{1}{R_0^2} \left(\frac{1}{\tilde{\lambda}} - 1 \right), \quad (63)$$

where

$$\tilde{\lambda} = \min_{\tilde{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \tilde{w}' z_t]^2, \quad (64)$$

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 (26), which is

a nonparametric estimator of $E[R_q^{+2}]$. However, when N is large and T is small, there is a high probability that we can find a portfolio that has negative gross return 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 $\hat{\sigma}_c^2$.¹⁵ To show this, we note that for any nonzero N -vector w_0 , we have

$$\begin{aligned}
 P[\hat{\lambda} = 0] &= P \left[\min_w \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 \mathbf{1}_N)]^2 = 0 \right] \\
 &\geq P \left[\sum_{t=1}^T \max[0, R_0 + w_0'(R_t - R_0 \mathbf{1}_N)]^2 = 0 \right] \\
 &= \prod_{t=1}^T P[R_0 + w_0'(R_t - R_0 \mathbf{1}_N) \leq 0] \\
 &= \Phi \left(-\frac{\mu_p}{\sigma_p} \right)^T, \tag{65}
 \end{aligned}$$

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.

4.4. 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

¹⁵Burnside (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.

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.,

$$P[\underline{\theta}_0^2 \leq \theta_0^2 \leq \bar{\theta}_0^2] = 1 - \alpha. \quad (66)$$

Then, using the fact that σ_0^2 and σ_c^2 are monotonically 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 .¹⁶ 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

$$F_{N,T-N,\bar{\theta}_0^2}(x) = \frac{\alpha}{2}, \quad (67)$$

$$F_{N,T-N,\underline{\theta}_0^2}(x) = 1 - \frac{\alpha}{2}, \quad (68)$$

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 decreasing in the noncentrality parameter δ , (68) 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 $\bar{\theta}_0^2 = 0$.¹⁷

¹⁶Lewellen, Nagel, and Shanken (2010) also use the statistical method to construct the confidence interval for the unexplained squared Sharpe ratio of an asset-pricing model.

¹⁷A Matlab program for computing the confidence intervals for σ_0^2 and σ_c^2 is available upon request.

Using the monthly excess returns (in excess of the one-month T-bill rate) on the 25 Fama-French size and book-to-market ranked portfolios (from Kenneth French’s website) over the period 1931/7–2014/10, the sample squared Sharpe ratio of the tangency portfolio is 0.0943 and the 95% confidence interval for the squared Sharpe ratio is (0.0361, 0.1081). Similar to the two HJ-bounds, the confidence interval for the squared Sharpe ratio is quite wide even for $T = 1000$, indicating that there is substantial uncertainty about the value of the squared Sharpe ratio of the tangency portfolio.

4.5. Effects of Nonnormality and Conditional Heteroskedasticity

The distributional results on the sample HJ-bounds in this paper are derived under the 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).

4.5.1. 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 to our finite-sample results that are derived under the normality assumption. Since we cannot derive the finite-sample distributions of $\hat{\sigma}_0^2$ and $\hat{\sigma}_c^2$ under the multivariate t -distribution assumption, we have to rely on simulation. 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, in Figure 12 we present 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 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 (the proof of this result is available upon request) that for the multivariate t -distribution with ν degrees of freedom, σ_c^2 is given by

$$\sigma_c^2 = \frac{1}{R_0^2} \left[\frac{\theta_0(\theta_0 + \eta^*)}{\Phi_{\nu-2}(\eta^*)} - 1 \right], \quad (69)$$

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

$$\frac{\eta \Phi_{\nu} \left(\left(\frac{\nu}{\nu-2} \right)^{\frac{1}{2}} \eta \right) + \phi_{\nu-2}(\eta)}{\Phi_{\nu-2}(\eta)} = \frac{1}{\theta_0}, \quad (70)$$

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.

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

4.5.2. 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 transformed excess returns, $z_t = P'V^{-\frac{1}{2}}r_t$, follow a constant correlation multivariate GARCH(1,1), 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 univariate GARCH(1,1) process

$$\begin{aligned}
 z_{it} &= E[z_{it}] + \epsilon_t, \\
 \epsilon_t &\sim N(0, h_t), \\
 h_t &= \omega_i + \alpha_i \epsilon_{t-1}^2 + \beta_i h_{t-1},
 \end{aligned} \tag{71}$$

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 . 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 ranked portfolios (from Kenneth French's website) over the post-World War II period (1946/1–2014/10). The average α_i and β_i parameter estimates across the 25 assets are 0.093 and 0.822, 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 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.

4.6. Sampling Distributions Using Historical Data

In this subsection, we compare the different estimators of the unconstrained and constrained HJ-bounds using real data. As in Section 4.4, we make use of the monthly returns on the 25 Fama-French size and book-to-market ranked portfolios (from Kenneth French’s website) over the period 1931/7–2014/10 ($T = 1000$). In Figure 16, we plot the estimated probability density functions (pdfs) of the sample HJ-bounds. For the cases in which the finite-sample distributions of the sample HJ-bounds are not available analytically, we run 100,000 simulations to characterize their finite-sample behavior.

Figure 16 about here

In the top panel of Figure 16, we plot the estimated pdfs of $\hat{\sigma}_0^2$ based on (1) its asymptotic distribution (Asymptotic) in (32). The pdf based on the asymptotic distribution derived under the normality assumption in (41) is not plotted in the graph since it is almost identical; (2) its finite-sample distribution under normality (Normal) in (48); and (3) its finite-sample distribution under stationary bootstrap resampling (Bootstrap).¹⁸ The figure shows that, even for $T = 1000$, the pdf of $\hat{\sigma}_0^2$ based on the asymptotic distribution is quite far from the pdfs based on the finite-sample distributions. In addition, the i.i.d. normal method based on (48) and the bootstrap resampling methods deliver estimated pdfs that are fairly close to each other.

Turning to the constrained HJ-bound, in the bottom panel of Figure 16, we plot the estimated pdfs of $\hat{\sigma}_c^2$ based on (1) its asymptotic distribution (Asymptotic) in (33). The pdf based on the asymptotic distribution derived under the normality assumption in (42) is not plotted in the graph since it is almost identical; (2) its finite-sample distribution under normality (Normal) that is obtained by simulation using Proposition 3; and (3) its finite-sample distribution under stationary bootstrap resampling (Bootstrap).¹⁹ For finite-sample comparison purposes, we also plot the estimated pdf of $\tilde{\sigma}_c^2$ based on its finite-sample distribution (Normal (MLE)) in (56). As for $\hat{\sigma}_0^2$, the asymptotic distribution of $\hat{\sigma}_c^2$ is quite far from its finite-sample counterparts even for $T = 1000$. In addition, the finite-sample distribution of $\tilde{\sigma}_c^2$ appears to be very close to the finite-sample distribution of $\hat{\sigma}_c^2$ under normality, and can be used as a reasonable approximation in finite samples. Finally, the normal pdf is much closer to the bootstrap pdf than the asymptotic one. Overall, our analysis suggests that the normality assumption, on which our proposed estimators and finite-sample results are built on, is not too restrictive even when real data is used to characterize the return distribution.

5. Conclusions

We provide a geometric interpretation of the unconstrained and constrained HJ-bounds in the mean-variance frontiers of returns and truncated returns, respectively. Assuming that the returns are multivariate normally distributed, we present the finite-sample distributions of the sample un-

¹⁸In implementing the stationary bootstrap of Politis and Romano (1994), the expected block length is set equal to 4. Given the little persistence in returns, different block lengths deliver results that are very similar.

¹⁹In computing the finite-sample distributions based on 100,000 simulations, we never encounter instances in which the sample constrained HJ-bound assumes the value of infinity.

constrained 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. We show the importance of our bias adjustment using simulations and an empirical application.

For typical number of assets and length of time series, the sample unconstrained and constrained 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. We believe it would be a good practice to report both the point estimates and the confidence intervals for the HJ-bounds in empirical work. Using simulations, 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 of a GARCH type.

Appendix

Proof of Lemma 1: (1) Since $E[R] = \mu$, the expectation of the second term in (7) vanishes and we have $E[m_0] = 1/R_0$. (2) Using the fact that $E[R(R - \mu)'] = V$, we have $E[m_0 R] = (\mu - (\mu - R_0 1_N))/R_0 = 1_N$. (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 portfolio with minimum second moment of its truncated return must satisfy the first order condition

$$E[R_{q^*}^+(R - R_0 1_N)] = 0_N. \quad (\text{A4})$$

Using (A4), we obtain

$$E[R_{q^*}^{+ 2}] = E[R_{q^*}^+ R_{q^*}] = E[R_{q^*}^+(R_0 + w^{*'}(R - R_0 1_N))] = R_0 E[R_{q^*}^+]. \quad (\text{A5})$$

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 $E[m_c R_0] = 1$, we obtain

$$E[m_c(R - R_0 1_N)] = 0_N \Rightarrow E[m_c R] = 1_N. \quad (\text{A6})$$

(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

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

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

This completes the proof.

Proof of Lemma 3: Using (A5), we can simplify the squared 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 moments of a truncated 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

$$\begin{aligned} E[R_p^+] &= \int_0^\infty R_p f(R_p) dR_p \\ &= \int_{-\eta}^\infty (\mu_p + \sigma_p y) \phi(y) dy \\ &= \mu_p \Phi(\eta) + \sigma_p \phi(\eta), \end{aligned} \tag{A17}$$

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

This completes the proof.

Proof of Lemma 5: Define

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

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

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

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

$$g'(u) = 1 + f'(u) = 1 - f(u) [u + f(u)] > 0$$

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 derivatives

$$\frac{\partial \mu_p}{\partial w} = \mu - R_0 1_N, \tag{A21}$$

$$\frac{\partial \sigma_p^2}{\partial w} = 2Vw, \tag{A22}$$

$$\frac{\partial \sigma_p}{\partial w} = \frac{1}{\sigma_p} Vw, \tag{A23}$$

$$\frac{\partial \Phi(\eta)}{\partial w} = \phi(\eta) \left[\frac{(\mu - R_0 1_N) \sigma_p - \eta Vw}{\sigma_p^2} \right], \tag{A24}$$

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

and Lemma 4, we can show that

$$\frac{\partial E[R_p^{+2}]}{\partial w} = \frac{\partial[(\mu_p^2 + \sigma_p^2)\Phi(\eta) + \mu_p\sigma_p\phi(\eta)]}{\partial w} = 2\Phi(\eta) [Vw + \sigma_p g(\eta)(\mu - R_0\mathbf{1}_N)]. \quad (\text{A26})$$

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

$$w^* = -\sigma_{q^*} g(\eta^*) V^{-1}(\mu - R_0\mathbf{1}_N). \quad (\text{A27})$$

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

$$\sigma_{q^*}^2 = w^{*'} V w^* = \sigma_{q^*}^2 g(\eta^*)^2 (\mu - R_0\mathbf{1}_N)' V^{-1} (\mu - R_0\mathbf{1}_N) = \sigma_{q^*}^2 g(\eta^*)^2 \theta_0^2. \quad (\text{A28})$$

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

$$g(\eta^*) = \frac{1}{\theta_0}, \quad (\text{A29})$$

and w^* must satisfy

$$w^* = -\frac{\sigma_{q^*}}{\theta_0} V^{-1}(\mu - R_0\mathbf{1}_N). \quad (\text{A30})$$

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

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

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

$$\sigma_{q^*} = \frac{R_0}{\eta^* + \theta_0}. \quad (\text{A32})$$

Substituting this into (A30), the portfolio with minimum second moment of its truncated return is given by

$$w^* = -\frac{R_0}{\theta_0(\eta^* + \theta_0)} V^{-1}(\mu - R_0\mathbf{1}_N). \quad (\text{A33})$$

Using (A5) and Lemma 3, we have

$$E[R_{q^*}^{+2}] = R_0 E[R_{q^*}^+] = R_0 [\mu_{q^*} \Phi(\eta^*) + \sigma_{q^*} \phi(\eta^*)] = R_0 \sigma_{q^*} \Phi(\eta^*) g(\eta^*) = \frac{R_0^2 \Phi(\eta^*)}{\theta_0(\eta^* + \theta_0)}, \quad (\text{A34})$$

where the last equality is obtained by using (A29) and (A32). From property (3) in Lemma 2, we have

$$E[m_c^2] = \frac{1}{E[R_{q^*}^{+2}]} = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)}. \quad (\text{A35})$$

Using property (1) in Lemma 2, we have $E[m_c] = 1/R_0 = \mu_m$. The expression for σ_c^2 can be obtained by using the fact that $\sigma_0 = \theta_0/R_0$. This completes the proof.

Proof of Lemma 6: We first derive $\partial\eta^*/\partial\theta_0$. Differentiating both sides of (16) with respect to η^* and using (A21) gives us

$$\begin{aligned} -\frac{1}{\theta_0^2} \frac{\partial\theta_0}{\partial\eta^*} &= 1 - \frac{\phi(\eta^*)}{\Phi(\eta^*)} \left[\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} \right] \\ \Rightarrow -\frac{1}{\theta_0^2} \frac{\partial\theta_0}{\partial\eta^*} &= 1 - \left(\frac{1}{\theta_0} - \eta^* \right) \frac{1}{\theta_0} \\ \Rightarrow \frac{\partial\eta^*}{\partial\theta_0} &= \frac{1}{1 - \theta_0(\eta^* + \theta_0)}. \end{aligned} \quad (\text{A36})$$

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

$$\frac{\partial\sigma_c^2}{\partial\theta_0} = \frac{\partial}{\partial\theta_0} \frac{\theta_0(\eta^* + \theta_0)}{R_0^2\Phi(\eta^*)} = \frac{\eta^* + \theta_0 \frac{\partial\eta^*}{\partial\theta_0} + 2\theta_0 - \theta_0(\eta^* + \theta_0) \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)}. \quad (\text{A37})$$

Using the fact that $\phi(\eta^*)/\Phi(\eta^*) = (1/\theta_0) - \eta^*$ and (A36), we can simplify the derivative to

$$\frac{\partial\sigma_c^2}{\partial\theta_0} = \frac{\eta^* + 2\theta_0 - \eta^*[1 - \theta_0(\eta^* + \theta_0)] \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)} = \frac{2\theta_0}{R_0^2\Phi(\eta^*)} > 0. \quad (\text{A38})$$

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

$$h'(\theta_0) = \frac{2\theta_0}{R_0^2\Phi(\eta^*)} - \frac{2\theta_0}{R_0^2} > 0 \quad (\text{A39})$$

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

$$\lim_{\theta_0 \rightarrow 0} \theta_0 \eta^* = \lim_{\theta_0 \rightarrow 0} \frac{\eta^*}{\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)}} = 1. \quad (\text{A40})$$

It follows that $\lim_{\theta_0 \rightarrow 0} \sigma_c^2 = \frac{1}{R_0^2} - \frac{1}{R_0^2} = 0$ and $\lim_{\theta_0 \rightarrow 0} h(\theta_0) = 0$. To prove 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

$$h''(\theta_0) = \frac{2 - 2\theta_0 \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)} - \frac{2}{R_0^2} = \frac{2[1 - \Phi(\eta^*)] - 2\theta_0 \frac{\phi(\eta^*)}{\Phi(\eta^*)} \frac{\partial\eta^*}{\partial\theta_0}}{R_0^2\Phi(\eta^*)}. \quad (\text{A41})$$

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

$$\begin{aligned}
\tilde{h}'(\theta_0) &= \frac{1}{\sigma_0^4} \left[\frac{2\theta_0}{R_0^2 \Phi(\eta^*)} \sigma_0^2 - \frac{2\theta_0}{R_0^2} \sigma_c^2 \right] \\
&= \frac{2\theta_0}{\sigma_0^4 R_0^4} \left[\frac{\theta_0^2}{\Phi(\eta^*)} - \frac{\theta_0(\eta^* + \theta_0)}{\Phi(\eta^*)} + 1 \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} \left[\frac{\Phi(\eta^*)}{\theta_0} - \eta^* \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} \left[\Phi(\eta^*) \left(\eta^* + \frac{\phi(\eta^*)}{\Phi(\eta^*)} \right) - \eta^* \right] \\
&= \frac{2\theta_0^2}{\sigma_0^4 R_0^4 \Phi(\eta^*)} [-\eta^* (1 - \Phi(\eta^*)) + \phi(\eta^*)]. \tag{A42}
\end{aligned}$$

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

$$u\Phi(u) + \phi(u) = \Phi(u) \left(u + \frac{\phi(u)}{\Phi(u)} \right), \tag{A43}$$

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

$$v_0 = E[\phi_{0,t}^2] = E \left[\left(\frac{(1 - u_t)^2}{R_0^2} - \frac{1 + \theta_0^2}{R_0^2} \right)^2 \right] = \frac{2\theta_0^2(2 + \theta_0^2)}{R_0^4}. \tag{A44}$$

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

$$E[m_{c,t}^2] = \frac{1}{\|R_{q^*,t}^+\|^2}, \tag{A45}$$

$$E[m_{c,t}^4] = \frac{E[(R_{q^*,t}^+)^4]}{\|R_{q^*,t}^+\|^8} = E[m_{c,t}^2]^4 E[(R_{q^*,t}^+)^4]. \tag{A46}$$

Using (A35), we have

$$E[m_{c,t}^2] = \frac{\theta_0(\eta^* + \theta_0)}{R_0^2 \Phi(\eta^*)}. \tag{A47}$$

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

$$\begin{aligned}
E[(R_p^+)^4] &= \int_{-\eta}^{\infty} (\mu_p + \sigma_p y)^4 \phi(y) dy \\
&= \mu_p^4 \Phi(\eta) + 4\mu_p^3 \sigma_p \phi(\eta) + 6\mu_p^2 \sigma_p^2 [-\eta \phi(\eta) + \Phi(\eta)] \\
&\quad + 4\mu_p \sigma_p^3 (\eta^2 + 2) \phi(\eta) + \sigma_p^4 [-(\eta^3 + 3\eta) \phi(\eta) + 3\Phi(\eta)] \\
&= [(\eta^4 + 6\eta^2 + 3)\Phi(\eta) + \eta(\eta^2 + 5)\phi(\eta)] \sigma_p^4
\end{aligned} \tag{A48}$$

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

$$\begin{aligned}
E[(R_{q^*,t}^+)^4] &= [(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)] \sigma_{q^*}^4 \\
&= \frac{[(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)] R_0^4}{(\eta^* + \theta_0)^4}
\end{aligned} \tag{A49}$$

by using the expression of σ_{q^*} in (A32). It follows that

$$E[m_{c,t}^4] = \frac{\theta_0^4 [(\eta^{*4} + 6\eta^{*2} + 3)\Phi(\eta^*) + \eta^*(\eta^{*2} + 5)\phi(\eta^*)]}{R_0^4 \Phi(\eta^*)^4}. \tag{A50}$$

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

$$\begin{aligned}
E[m_{c,t}^4] &= \frac{\theta_0^4 \left[(\eta^{*4} + 6\eta^{*2} + 3) + \eta^*(\eta^{*2} + 5) \left(\frac{1}{\theta_0} - \eta^* \right) \right]}{R_0^4 \Phi(\eta^*)^3} \\
&= \frac{\theta_0^4 (\eta^{*2} + 3)}{R_0^4 \Phi(\eta^*)^3} + \frac{\eta^* \theta_0^3 (\eta^{*2} + 5)}{R_0^4 \Phi(\eta^*)^3} \\
&= \frac{\theta_0^3 [\theta_0 (\eta^{*2} + 3) + \eta^* (\eta^{*2} + 5)]}{R_0^4 \Phi(\eta^*)^3}.
\end{aligned} \tag{A51}$$

It follows that

$$\begin{aligned}
E[\phi_{c,t}^2] &= E[m_{c,t}^4] - 2(\sigma_c^2 + \mu_m^2) E[m_{c,t}^2] + (\sigma_c^2 + \mu_m^2)^2 \\
&= E[m_{c,t}^4] - E[m_{c,t}^2]^2 \\
&= \frac{\theta_0^3 [\theta_0 (\eta^{*2} + 3) + \eta^* (\eta^{*2} + 5)]}{R_0^4 \Phi(\eta^*)^3} - \frac{\theta_0^2 (\eta^* + \theta_0)^2}{R_0^4 \Phi(\eta^*)^2}.
\end{aligned} \tag{A52}$$

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 (A38), we obtain $\partial \sigma_c^2 / \partial \sigma_0^2 = 1 / \Phi(\eta^*)$. Using the delta method, we obtain

$$\sqrt{T}(\tilde{\sigma}_c^2 - \sigma_c^2) \overset{A}{\rightsquigarrow} N\left(0, \frac{v_0}{\Phi(\eta^*)^2}\right). \tag{A53}$$

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

$$\hat{\mu} \sim N(\mu, V/T), \quad (\text{A54})$$

$$T\hat{V} \sim W_N(T-1, V), \quad (\text{A55})$$

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

$$Y \equiv \frac{T\tilde{\theta}_0^2}{\hat{\theta}_0^2} \sim \chi_{T-N}^2, \quad (\text{A56})$$

and the ratio is independent of $\hat{\mu}$ and hence independent of $\tilde{\theta}_0^2$. From (A54), we have

$$\hat{\mu} - R_0 \mathbf{1}_N \sim N(\mu - R_0 \mathbf{1}_N, V/T). \quad (\text{A57})$$

Therefore,

$$X \equiv T\tilde{\theta}_0^2 \sim \chi_N^2(T\theta_0^2) \quad (\text{A58})$$

and is independent of Y . Together, we have

$$\hat{\theta}_0^2 = \frac{T\tilde{\theta}_0^2}{Y} = \frac{X}{Y} \sim \left(\frac{N}{T-N} \right) F_{N, T-N}(T\theta_0^2). \quad (\text{A59})$$

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

$$E[\hat{\theta}_0^r] = E[X^{\frac{r}{2}}] E[Y^{-\frac{r}{2}}]. \quad (\text{A60})$$

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

$$E[X^{\frac{r}{2}}] = \frac{2^{\frac{r}{2}} \Gamma(\frac{N+r}{2})}{\Gamma(\frac{N}{2})} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right), \quad (\text{A61})$$

$$E[Y^{-\frac{r}{2}}] = \frac{2^{-\frac{r}{2}} \Gamma(\frac{T-N-r}{2})}{\Gamma(\frac{T-N}{2})}. \quad (\text{A62})$$

Combining these two expressions, we obtain

$$E[\hat{\theta}_0^r] = \frac{\Gamma(\frac{N+r}{2}) \Gamma(\frac{T-N-r}{2})}{\Gamma(\frac{N}{2}) \Gamma(\frac{T-N}{2})} {}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right). \quad (\text{A63})$$

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

$${}_1F_1\left(-\frac{r}{2}; \frac{N}{2}; -\frac{T\theta_0^2}{2}\right) = \sum_{i=0}^{r/2} \frac{\left(-\frac{r}{2}\right)_i - \left(\frac{T\theta_0^2}{2}\right)^i}{\left(\frac{N}{2}\right)_i i!} = \sum_{i=0}^{r/2} \frac{\binom{r/2}{i} \left(\frac{T\theta_0^2}{2}\right)^i}{\left(\frac{N}{2}\right)_i}. \quad (\text{A64})$$

This completes the proof.

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

$$\hat{\sigma}_c^2 = \frac{1}{R_0^2} \left(\frac{1}{\tilde{\lambda}} - 1 \right), \quad (\text{A65})$$

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

$$\tilde{\lambda} = \frac{1}{R_0^2} \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'(R_t - R_0 \mathbf{1}_N)]^2. \quad (\text{A66})$$

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

$$\tilde{\lambda} = \frac{1}{R_0^2} \min_w \frac{1}{T} \sum_{t=1}^T \max[0, R_0 + w'r_t]^2 = \min_{\hat{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \hat{w}'r_t]^2, \quad (\text{A67})$$

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

$$z_t \equiv P'V^{-\frac{1}{2}}r_t \sim N([\theta_0, \mathbf{0}'_{N-1}]', I_N). \quad (\text{A68})$$

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

$$\tilde{\lambda} = \min_{\tilde{w}} \frac{1}{T} \sum_{t=1}^T \max[0, 1 + \tilde{w}'z_t]^2, \quad (\text{A69})$$

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

References

- Aldrich, E. M., A. R. Gallant. 2011. Habit, long-run risks, prospect? A statistical inquiry. *Journal of Financial Econometrics* **9**(4) 589–618.
- Backus, D. K., A. W. Gregory, C. I. Telmer. 1993. Accounting for forward rates in markets for foreign currency. *Journal of Finance* **48**(5) 1887–1908.
- Bansal, R., D. Kiku, A. Yaron. 2012a. An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review* **1**(1) 183–221.
- Bansal, R., D. Kiku, A. Yaron. 2012b. Risks for the long run: Estimation with time aggregation. *NBER* working paper 18305, Boston, MA.
- Bekaert, G., R. J. Hodrick. 1992. Characterizing predictable components in excess returns on equity and in foreign exchange markets. *Journal of Finance* **47**(2) 467–510.
- Bekaert, G., J. Liu. 2004. Conditioning information and variance bounds on pricing kernels. *Review of Financial Studies* **17**(2) 339–378.
- Birnbaum, Z. W. 1942. An inequality for Mill’s ratio. *Annals of Mathematical Statistics* **13**(2) 245–246.
- Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **31**(3) 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**(3) 498–505.
- Britten-Jones, M. 1999. The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance* **54**(2) 655–671.
- Burnside, A. C. 1994. Hansen-Jagannathan bounds as classical tests of asset pricing models. *Journal of Business and Economic Statistics* **12**(1) 57–79.
- Campbell, J. Y., J. Cochrane. 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* **107**(2) 205–251.

- Casella, G., R. L. Berger. 1990. *Statistical Inference*. Duxbury Press, CA, USA.
- Cecchetti, S. G., P. Lam, N. C. Mark. 1992. Testing volatility restrictions on intertemporal marginal rates of substitution implied by Euler equations and asset returns. *NBER technical paper 124*, Boston, MA.
- Cecchetti, S. G., P. Lam, N. C. Mark. 1994. Testing volatility restrictions on intertemporal marginal rates of substitution implied by Euler equations and asset returns. *Journal of Finance* **49**(1) 123–152.
- Cochrane, J. H., L. P. Hansen. 1992. Asset pricing explorations for macroeconomics. *NBER Macroeconomics Annual* **7** 115–165.
- Epstein, L. G., S. E. Zin. 2001. The independence axiom and asset returns. *Journal of Empirical Finance* **8**(5) 537–572.
- Favero, C. A., F. Ortu, A. Tamoni, H. Yang. 2014. Implications of returns predictability across horizons for asset pricing models. Working paper, Bocconi University, Milan, Italy.
- Ferson, W. E., A. F. Siegel. 2003. Stochastic discount factor bounds with conditioning information. *Review of Financial Studies* **16**(2) 567–595.
- Gibbons, M. R., S. A. Ross, J. Shanken. 1989. A test of the efficiency of a given portfolio. *Econometrica* **57**(5) 1121–1152.
- Gordon, R. D. 1941. Values of Mills’ ratio of area to bounding ordinate and of the normal probability integral for large values of the argument. *Annals of Mathematical Statistics* **12**(3) 364–366.
- Gordon, S., L. Samson, B. Carmichael. 1995. Finite-sample inferences about mean-standard deviation bounds for stochastic discount factors. *Economics Letters* **49**(3) 295–300.
- Gregory, A. W., G. W. Smith. 1992. Sampling variability in Hansen-Jagannathan bounds. *Economics Letters* **38**(3) 263–267.
- Hagiwara, M., M. A. Herce. 1997. Risk aversion and stock price sensitivity to dividends. *American Economic Review* **87**(4) 738–745.

- Hansen, L. P., J. Heaton, E. G. J. Luttmer. 1995. Econometric evaluation of asset pricing models. *Review of Financial Studies* **8**(2) 237–274.
- Hansen, L. P., R. Jagannathan. 1991. Implications of security market data for models of dynamic economies. *Journal of Political Economy* **99**(2) 225–262.
- He, H., D. M. Modest. 1995. Market frictions and consumption-based asset pricing. *Journal of Political Economy* **103**(1) 94–117.
- Heaton, J. 1995. An empirical investigation of asset pricing with temporally dependent preference specifications. *Econometrica* **63**(3) 681–717.
- Johnson, N. L., S. Kotz, N. Balakrishnan. 1995. *Continuous Univariate Distributions* **2**. Wiley, NY, USA.
- Kan, R., G. Zhou. 2006. A new variance bound on the stochastic discount factor. *Journal of Business* **79**(2) 941–961.
- Lewellen, J. W., S. Nagel, J. Shanken. 2010. A skeptical appraisal of asset pricing tests. *Journal of Financial Economics* **96**(2) 175–194.
- Muirhead, R. J. 1982. *Aspects of Multivariate Statistical Theory*. Wiley, NY, USA.
- Nakamura, E., J. Steinsson, R. Barro, J. Ursúa. 2013. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics* **5**(3) 35–74.
- Otrok, C., B. Ravikumar, C. H. Whiteman. 2002. Evaluating asset-pricing models using the Hansen-Jagannathan bound: A Monte Carlo investigation. *Journal of Applied Econometrics* **17**(2) 149–174.
- Politis, D. N., J. P. Romano. 1994. The stationary bootstrap. *Journal of the American Statistical Association* **89** 1303–1313.
- Sampford, M. R. 1953. Some inequalities on Mill’s ratio and related functions. *Annals of Mathematical Statistics* **24**(1) 130–132.
- Snow, K. N. 1991. Diagnosing asset pricing models using the distribution of asset returns. *Journal of Finance* **46**(3) 955–984.




Figure 1 Sample Hansen-Jagannathan Bounds. The figure presents the sample unconstrained Hansen-Jagannathan bound (solid line) and the nonparametric estimator of the constrained Hansen-Jagannathan bound (dashed line). The sample Hansen-Jagannathan bounds are computed using annual real non-overlapping returns on four size and book-to-market ranked portfolios, five industry portfolios, and a three-month Treasury bill over the period 1952–2012.

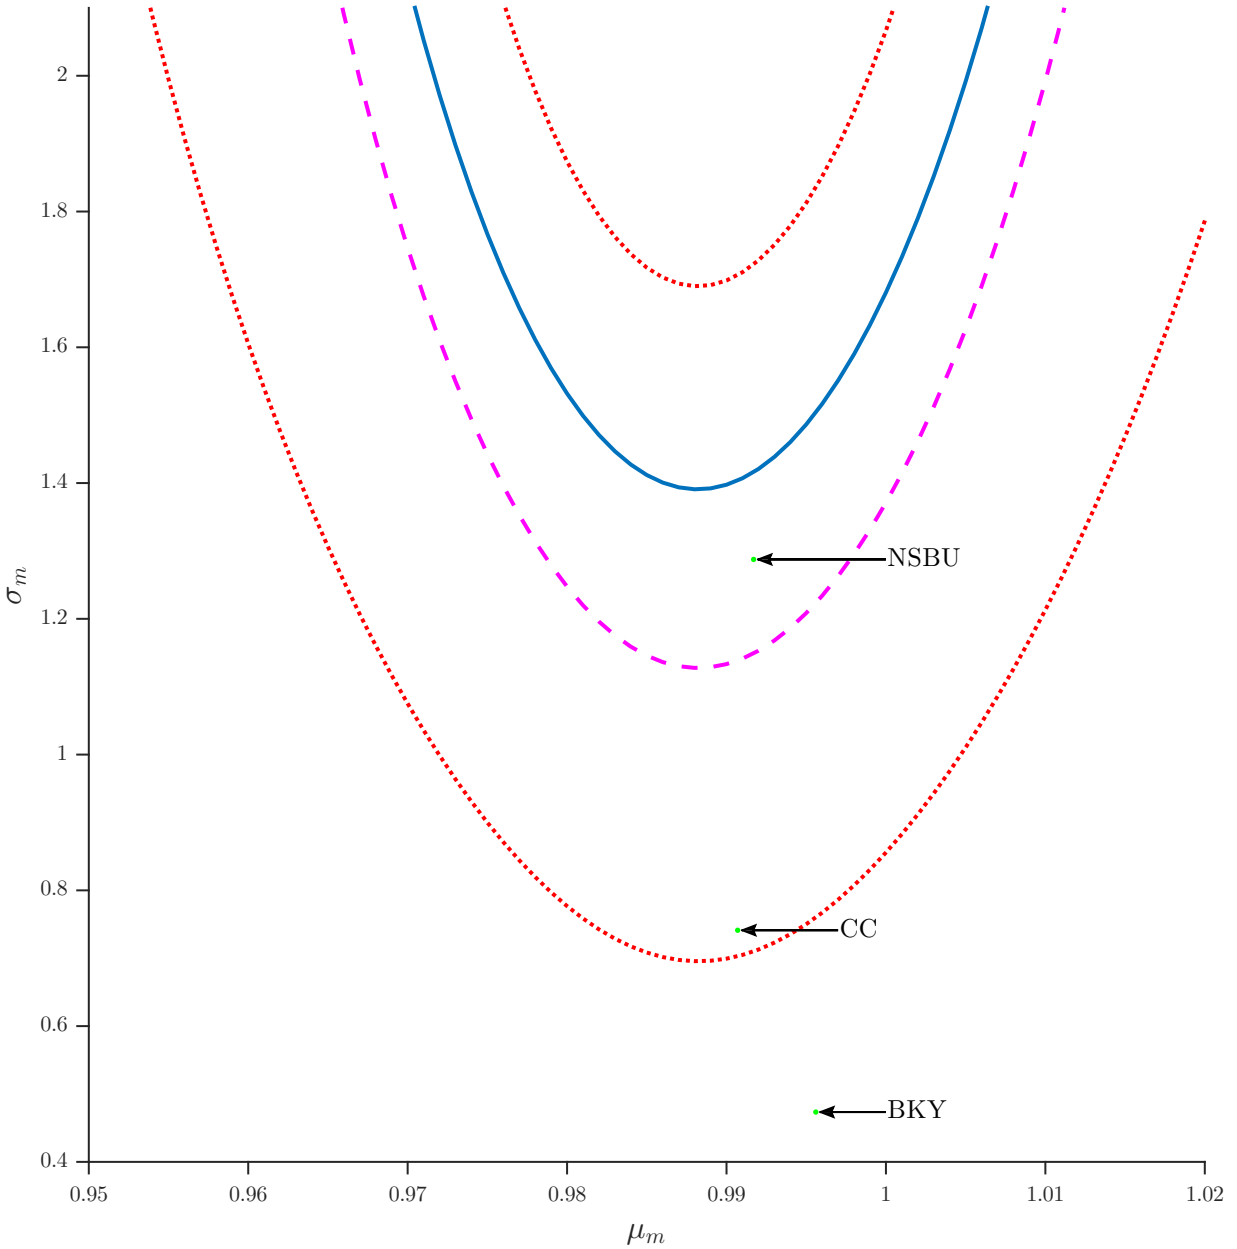


Figure 2 Confidence Intervals for the Constrained Hansen-Jagannathan Bound. The figure presents the mean-standard deviation pairs of the stochastic discount factors implied by the long-run risks model of Bansal, Kiku, and Yaron (BKY, 2012a), the external habit model of Campbell and Cochrane (CC, 1999), and the rare disasters model of Nakamura, Steinsson, Barro, and Ursúa (NSBU, 2013). The figure also presents the maximum likelihood estimator (solid line) and an approximate unbiased estimator (dashed line) 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 non-overlapping returns on four size and book-to-market ranked portfolios, five industry portfolios, and a three-month Treasury bill over the period 1952–2012.




Figure 3 Graphical Representation of the Unconstrained Hansen-Jagannathan Bound. The figure provides a geometric interpretation of the unconstrained HJ-bound using the mean-variance frontier of the 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 (σ_p, μ_p) space. 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 geometric interpretation of the constrained HJ-bound using the mean-variance frontier of the truncated portfolio returns. R_0 is the gross risk-free rate. The two curves emanating from the point $(0, R_0)$ represent the minimum/maximum variance frontier of the truncated portfolio returns in the (σ_p^+, μ_p^+) space. The portfolio with minimum second moment of its truncated return is represented by q^* , and it is the portfolio that is closest to the origin. The absolute value of the 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 $E[m] = 1/R_0$.

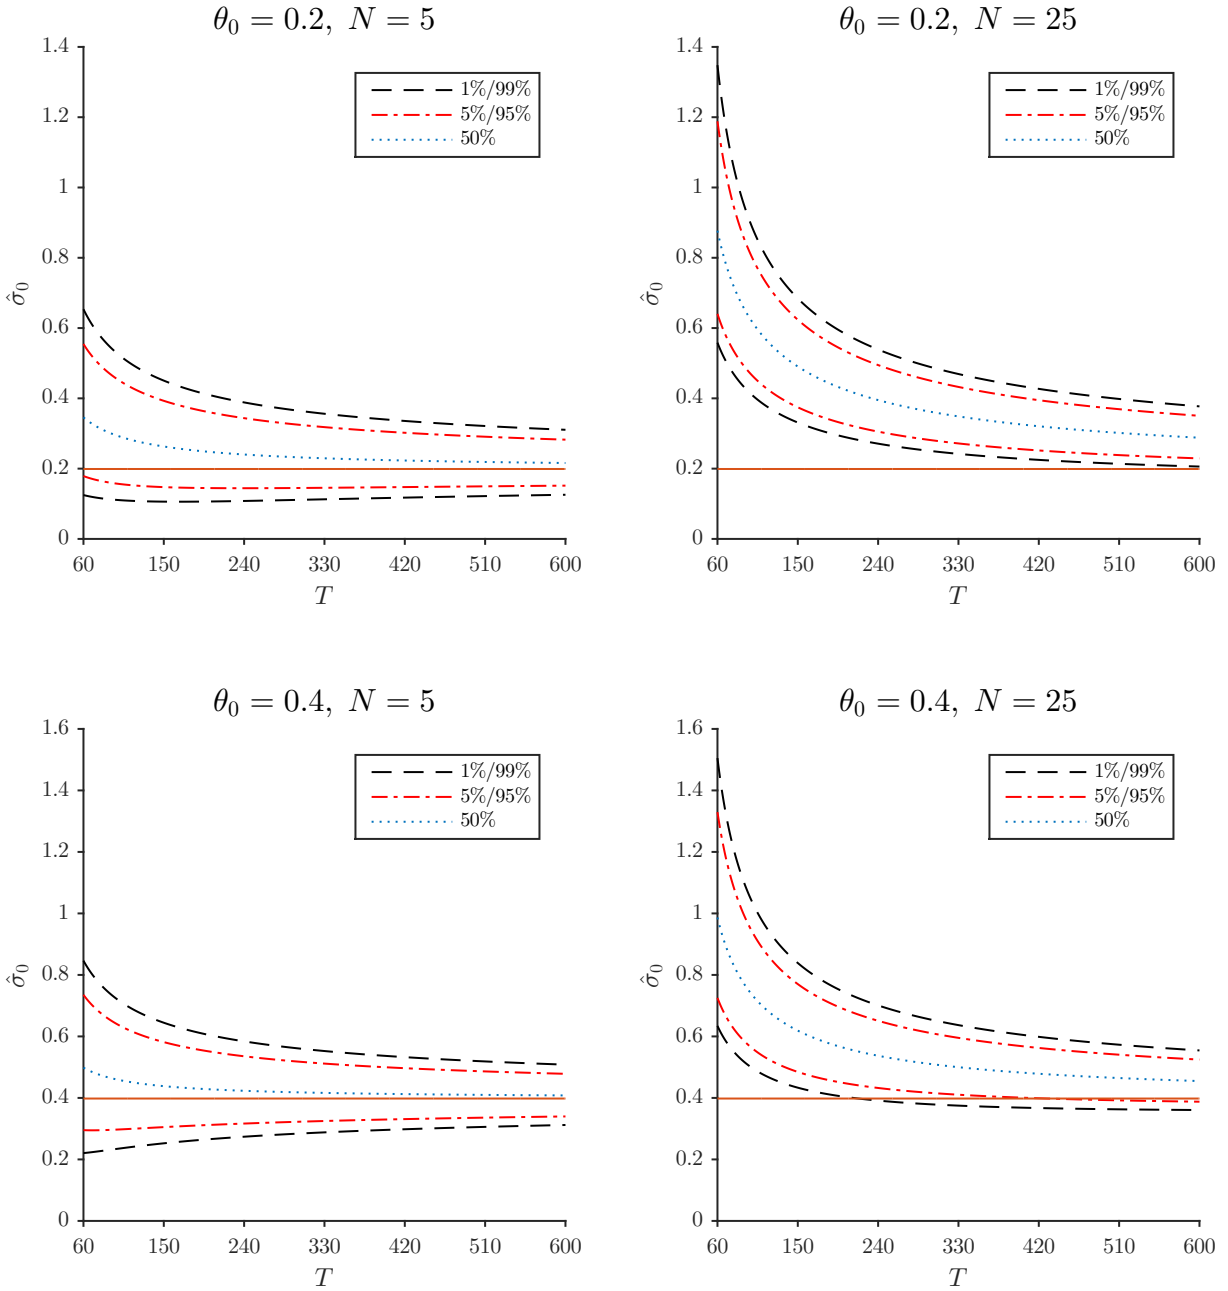


Figure 5 Exact Distribution of the Sample Unconstrained Hansen-Jagannathan Bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the sample unconstrained HJ-bound for different number of risky assets (N) and length of time series observations (T) under 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.

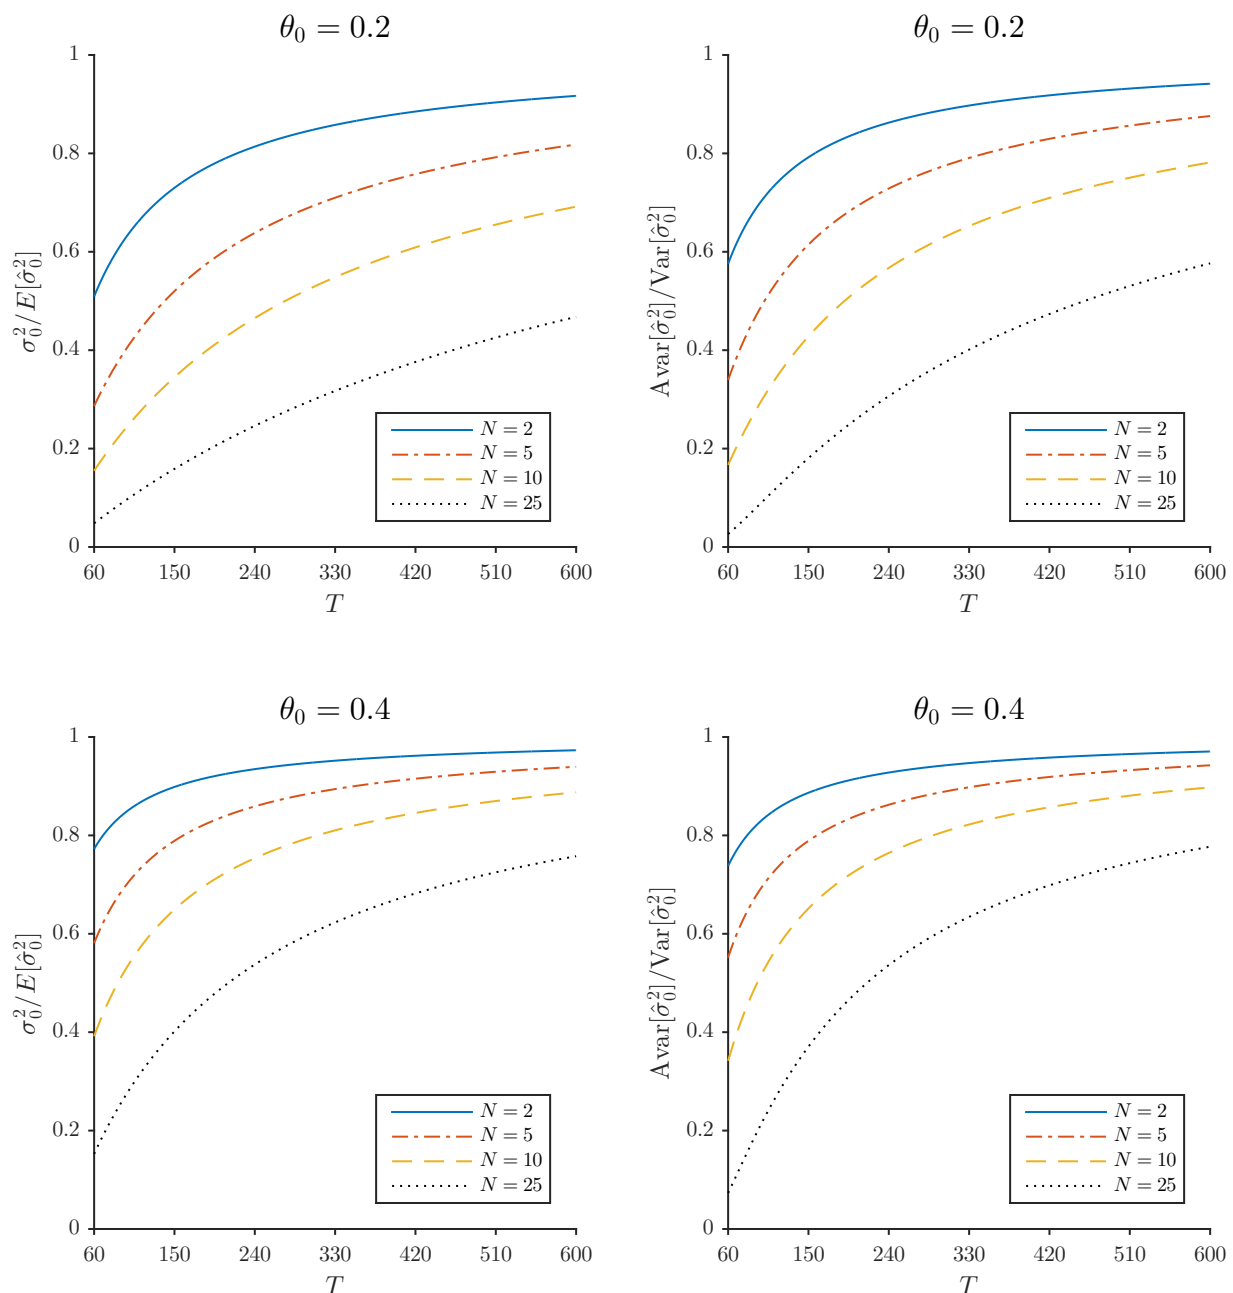


Figure 6 Ratio of Asymptotic to Exact Moments of the Sample Unconstrained Hansen-Jagannathan Bound. The figure presents the ratios of the asymptotic mean and variance to the exact mean and variance of the sample unconstrained HJ-bound ($\hat{\sigma}_0^2$) for different number of risky assets (N) and 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.

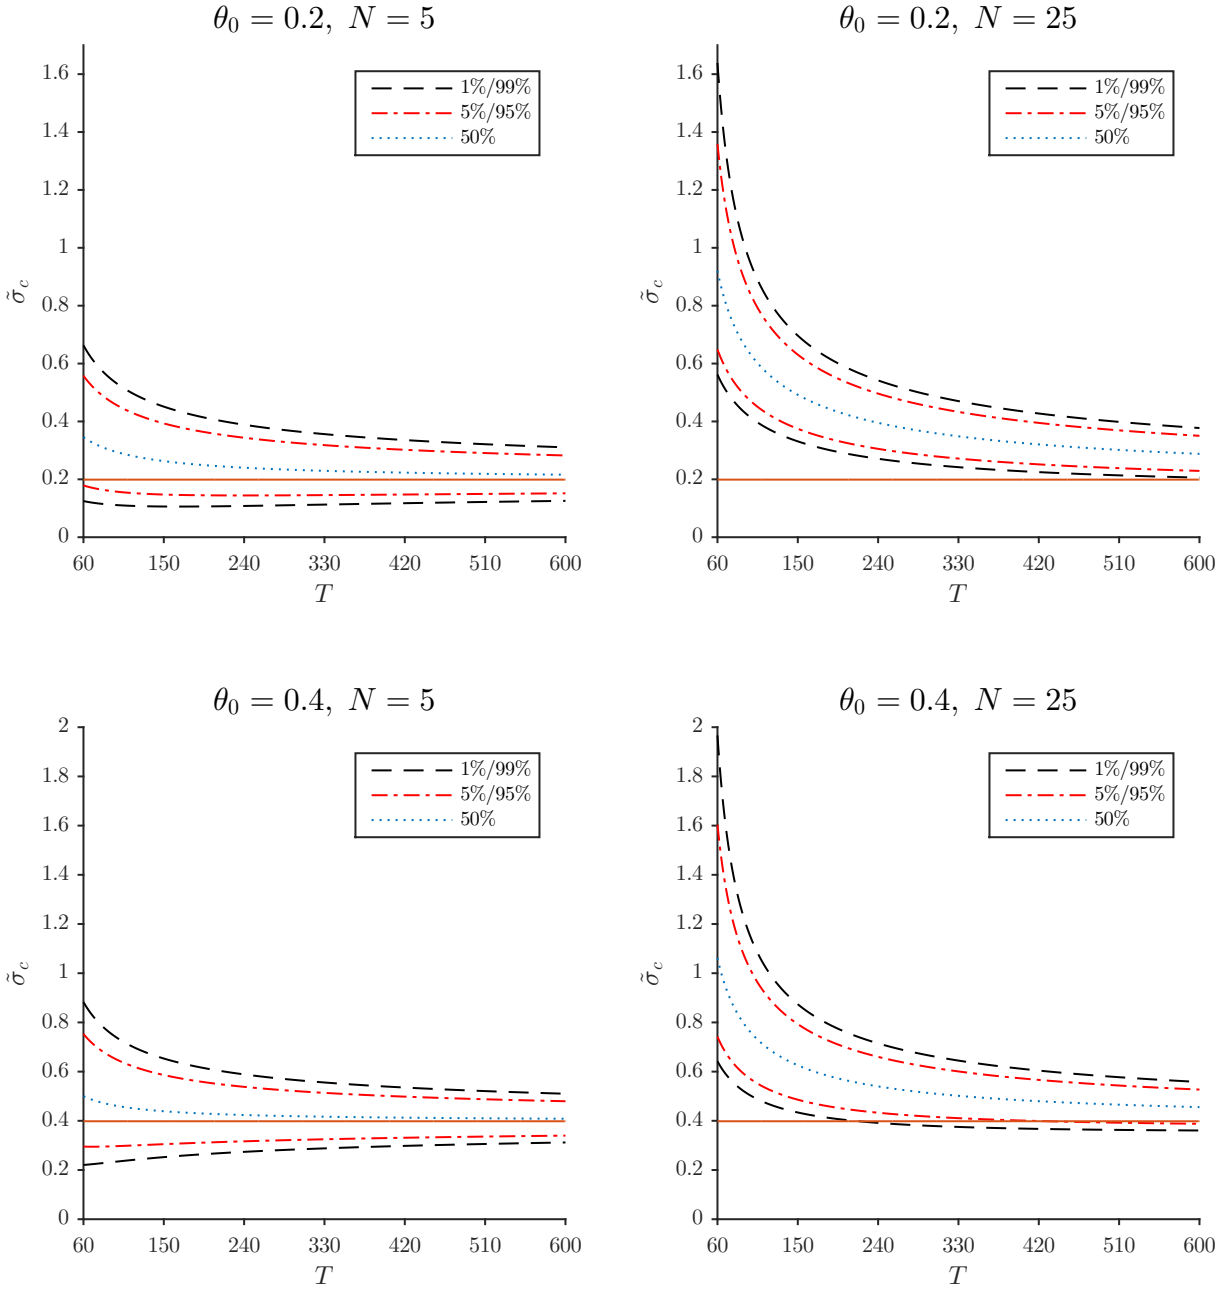


Figure 7 Exact Distribution of the Maximum Likelihood Estimator of the Constrained Hansen-Jagannathan Bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the maximum likelihood estimator of the constrained HJ-bound for different number of risky assets (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.

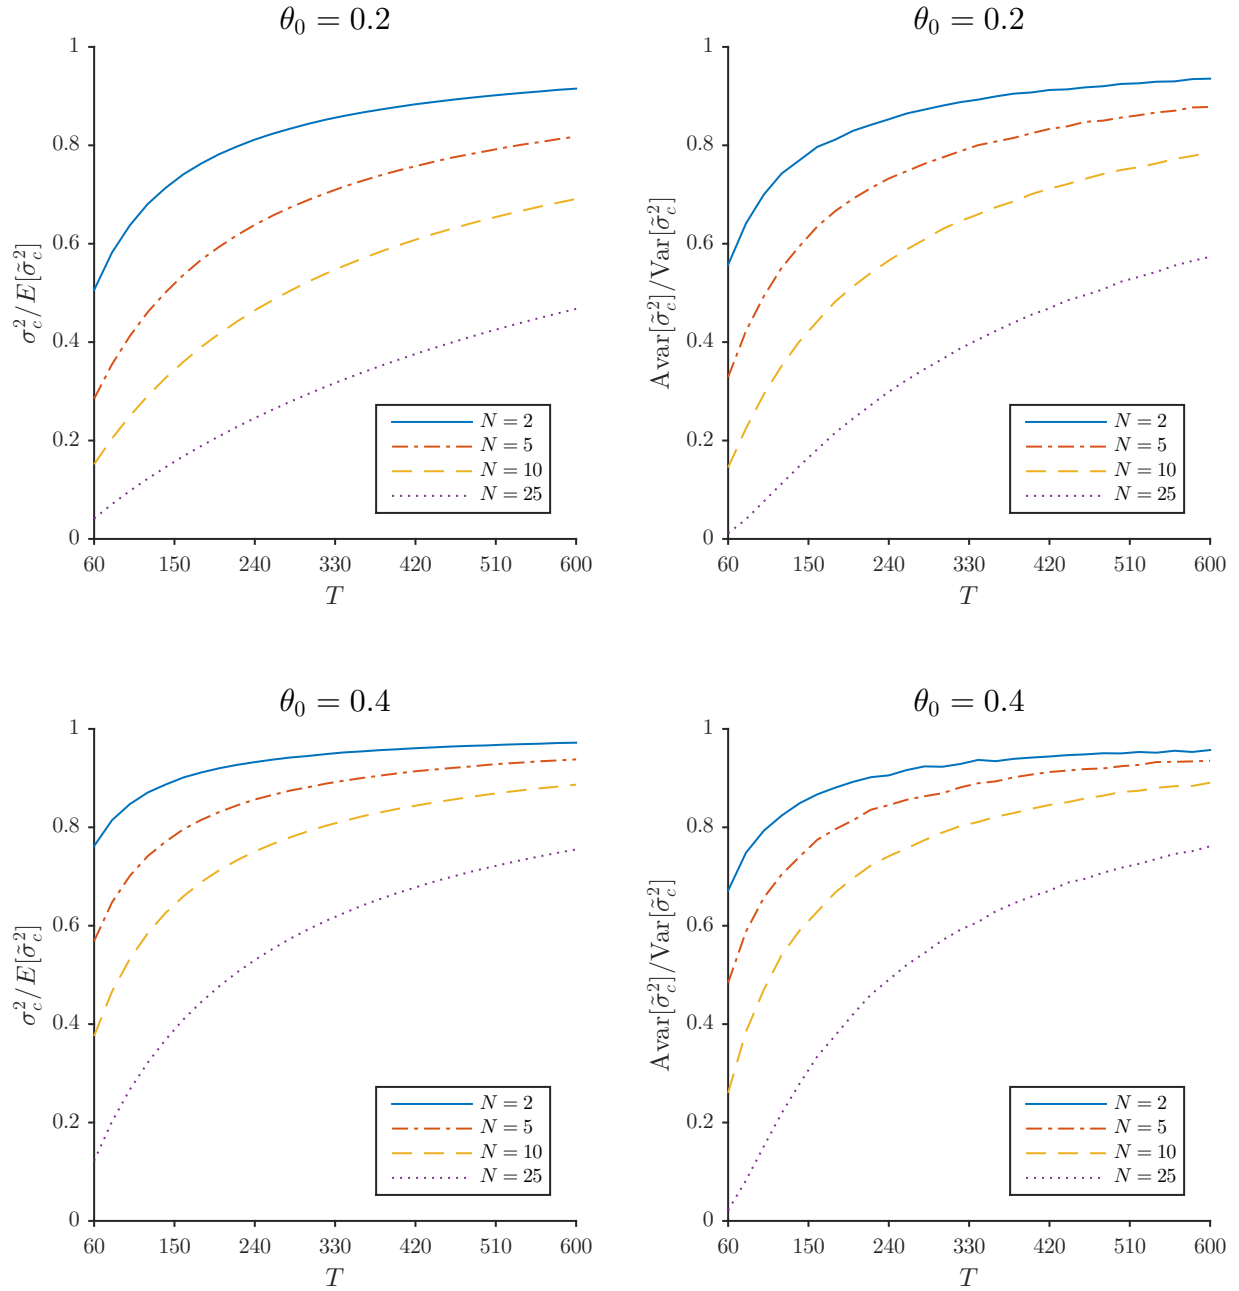


Figure 8 Ratio of Asymptotic to Exact Moments of the Maximum Likelihood Estimator of the Constrained Hansen-Jagannathan Bound. The figure presents the ratios of the asymptotic mean and variance to the exact mean and variance of the maximum likelihood estimator of the constrained HJ-bound ($\tilde{\sigma}_c^2$) 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.

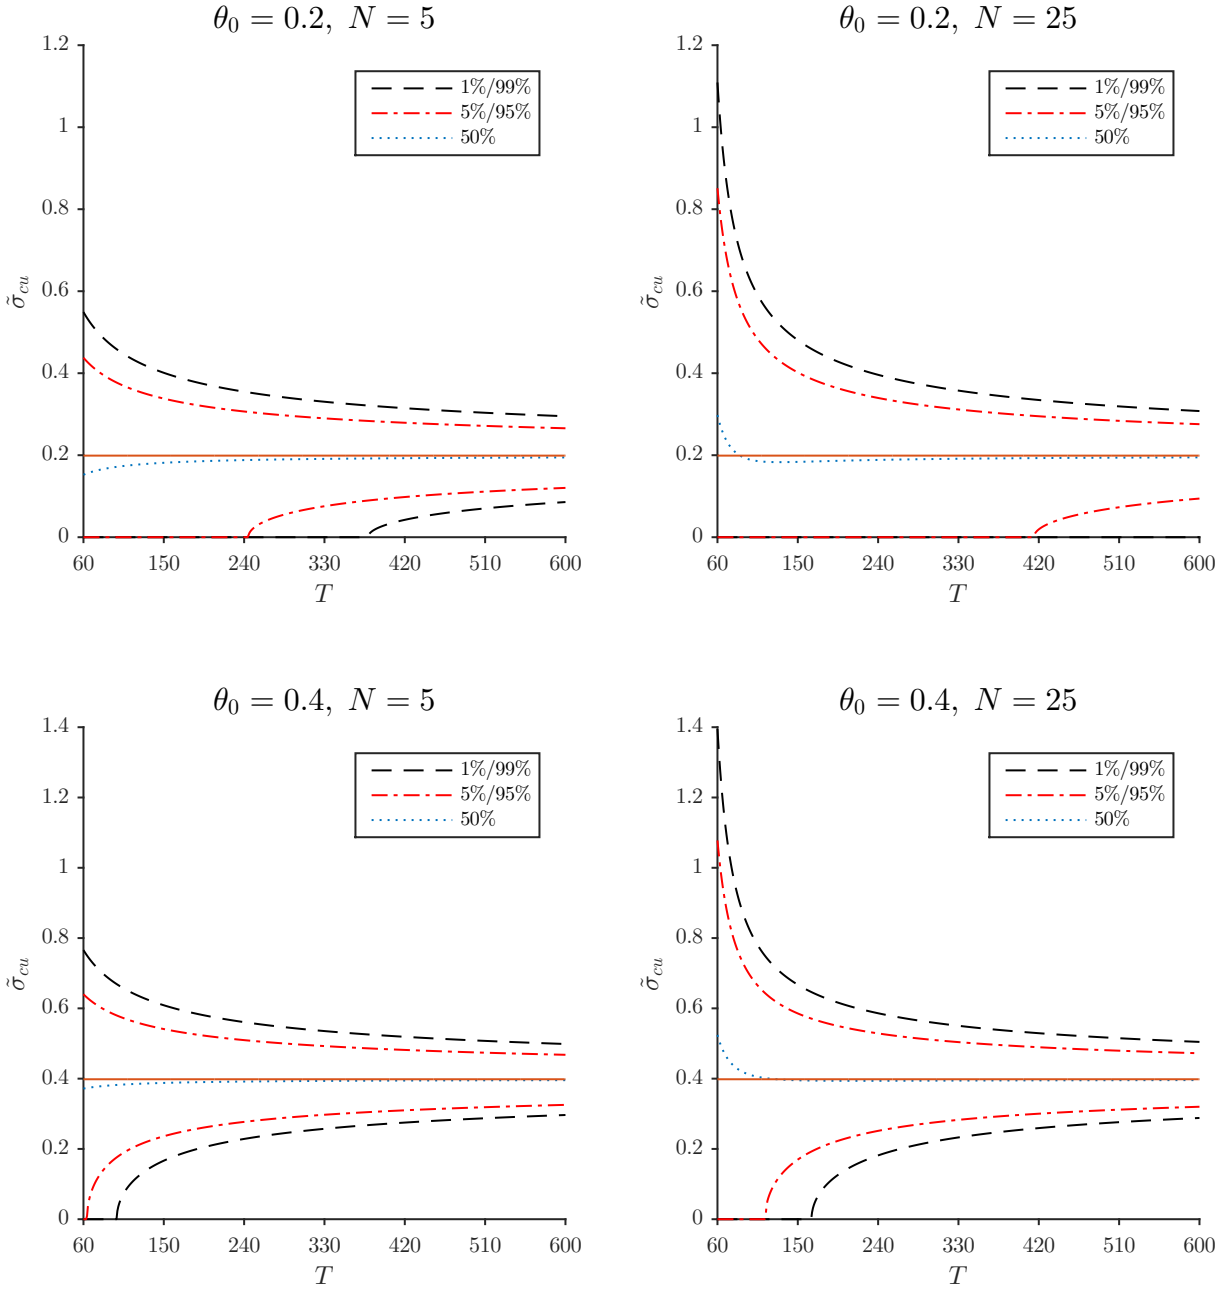


Figure 9 Finite-Sample Distribution of an Approximate Unbiased Estimator of the Constrained Hansen-Jagannathan Bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the finite-sample distribution of an approximate unbiased estimator of the constrained HJ-bound for different number 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.

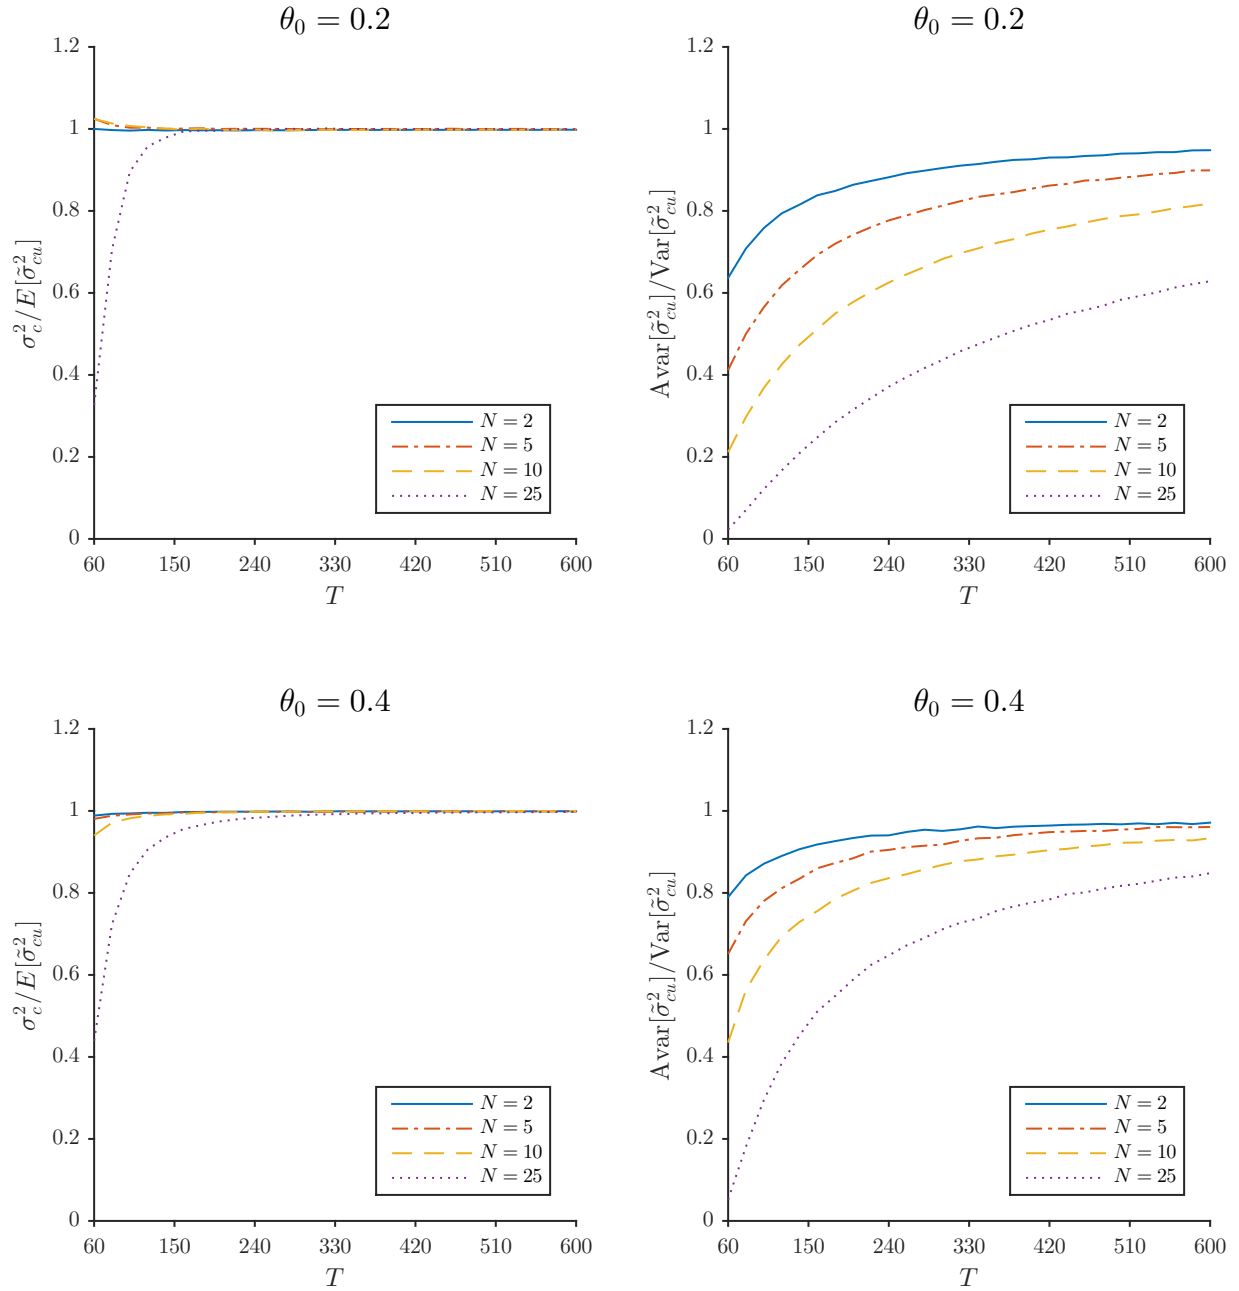


Figure 10 Ratio of Asymptotic to Finite Moments of an Approximate Unbiased Estimator of the Constrained Hansen-Jagannathan Bound. The figure presents the ratios of the asymptotic mean and variance to the finite mean and variance of an approximate unbiased estimator of the constrained HJ-bound ($\tilde{\sigma}_{cu}^2$) 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.

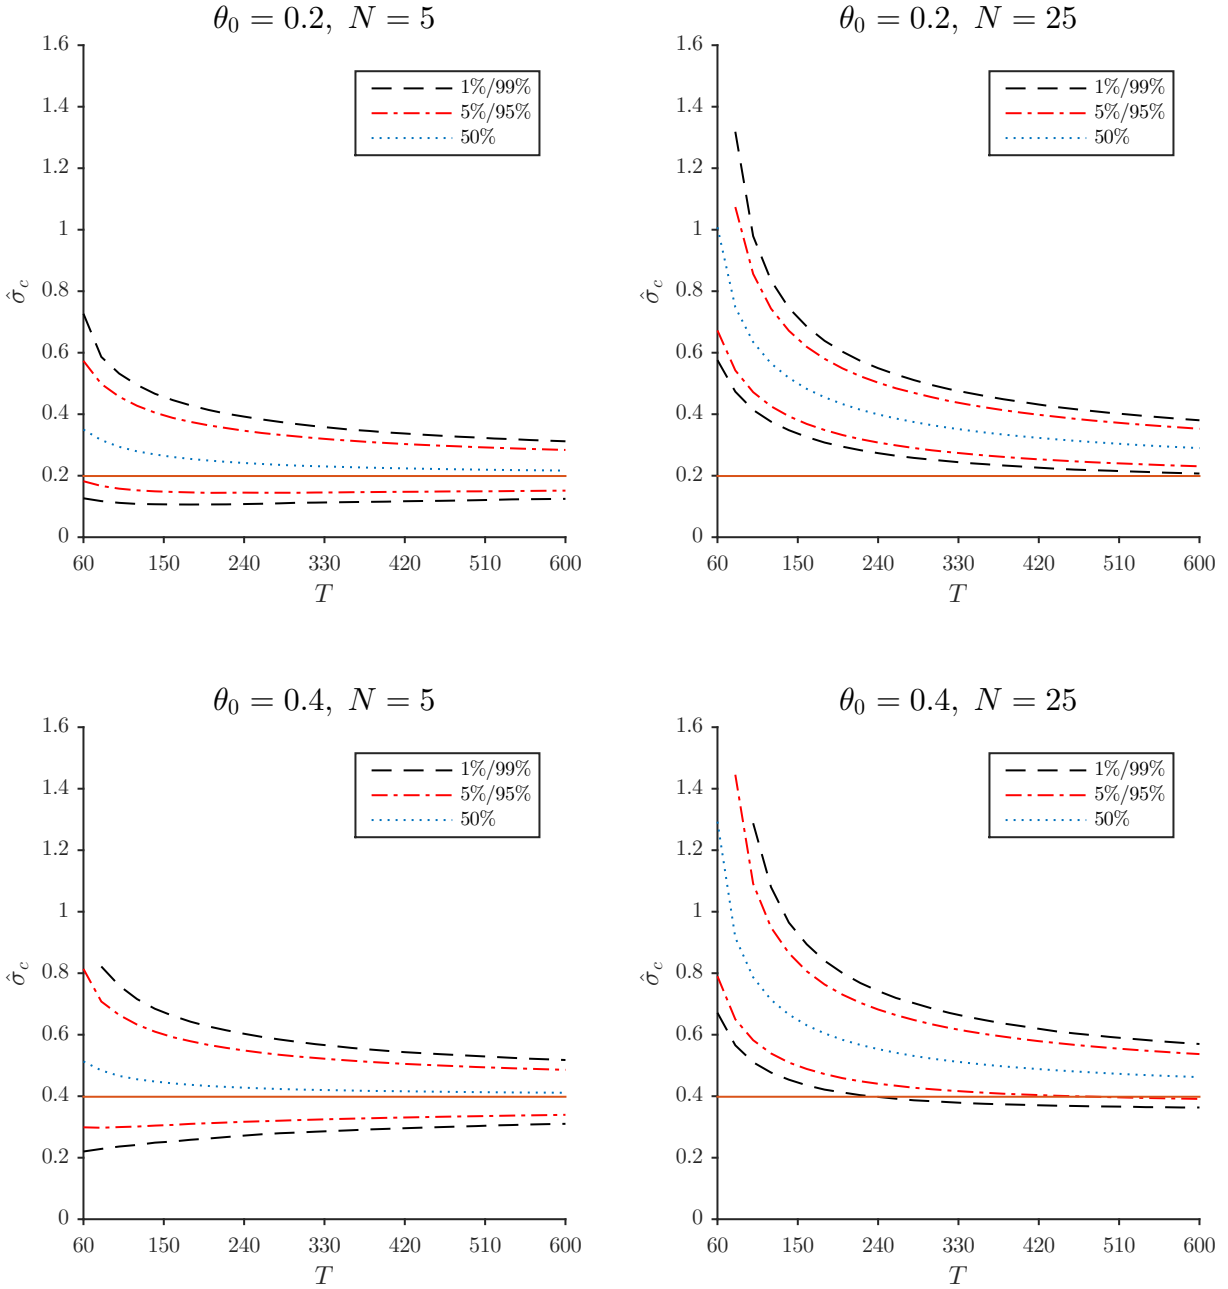


Figure 11 Exact Distribution of the Nonparametric Estimator of the Constrained Hansen-Jagannathan Bound. The figure presents the 1st, 5th, 50th, 95th, and 99th percentiles of the exact distribution of the estimator of the constrained HJ-bound for different number of risky assets (N) and length of time 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.

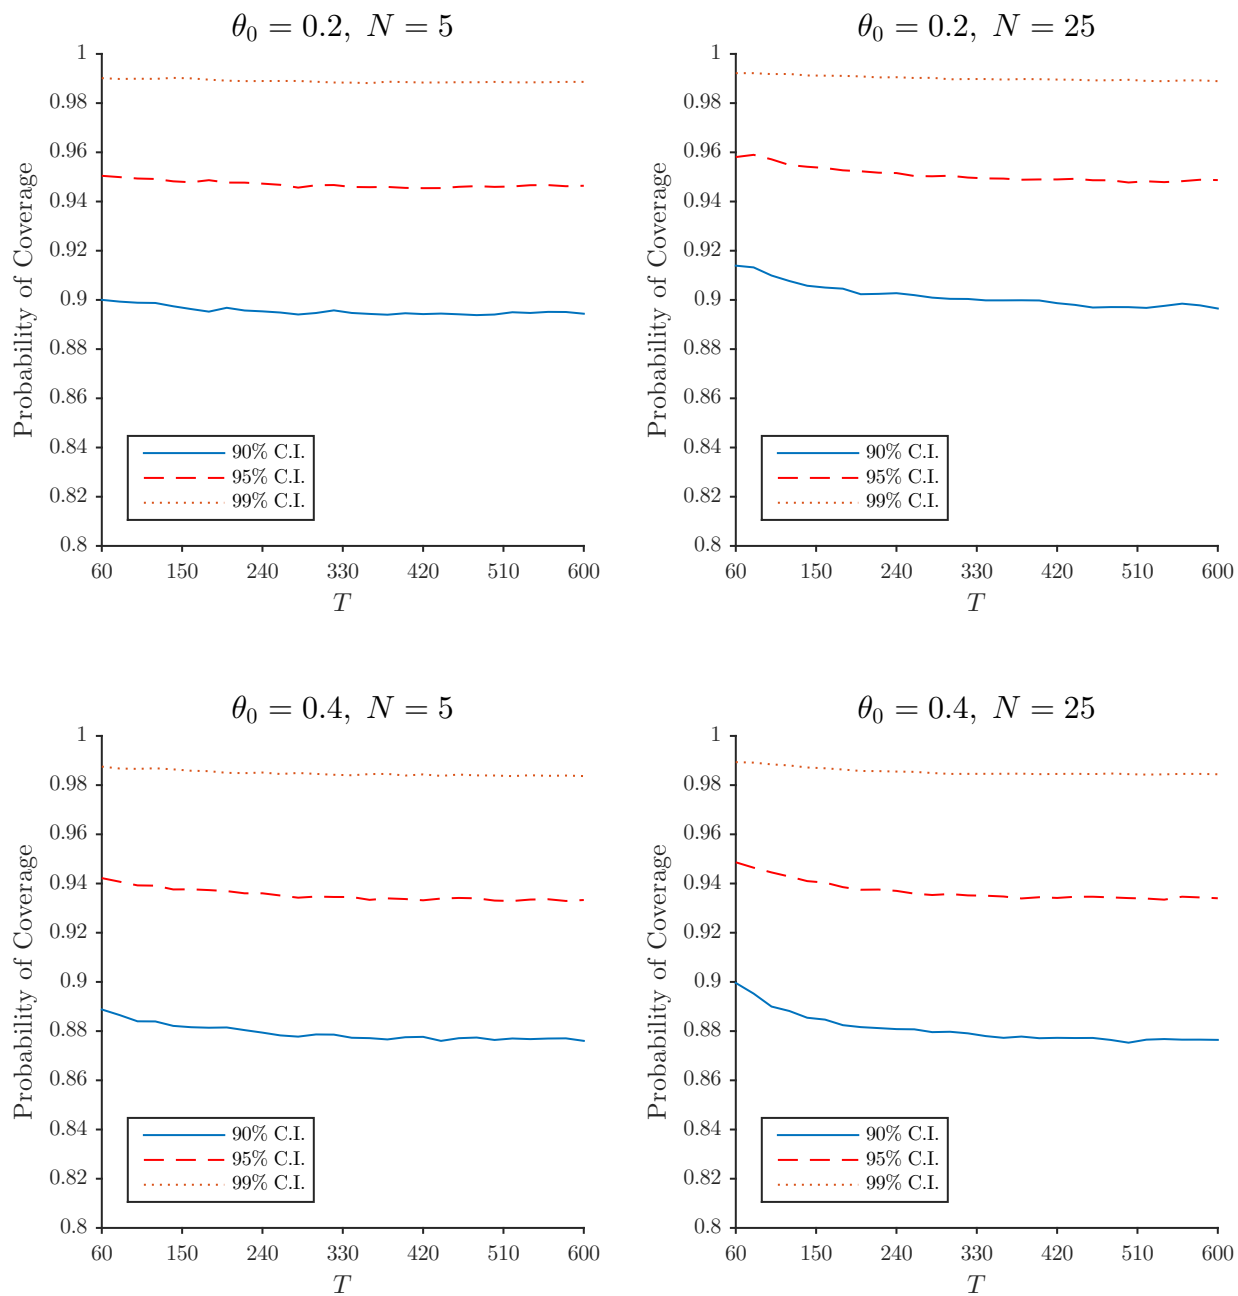


Figure 12 Probabilities of Coverage of the Confidence Intervals for the Unconstrained Hansen-Jagannathan Bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the unconstrained HJ-bound for different number of risky assets (N) and length of time series 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.

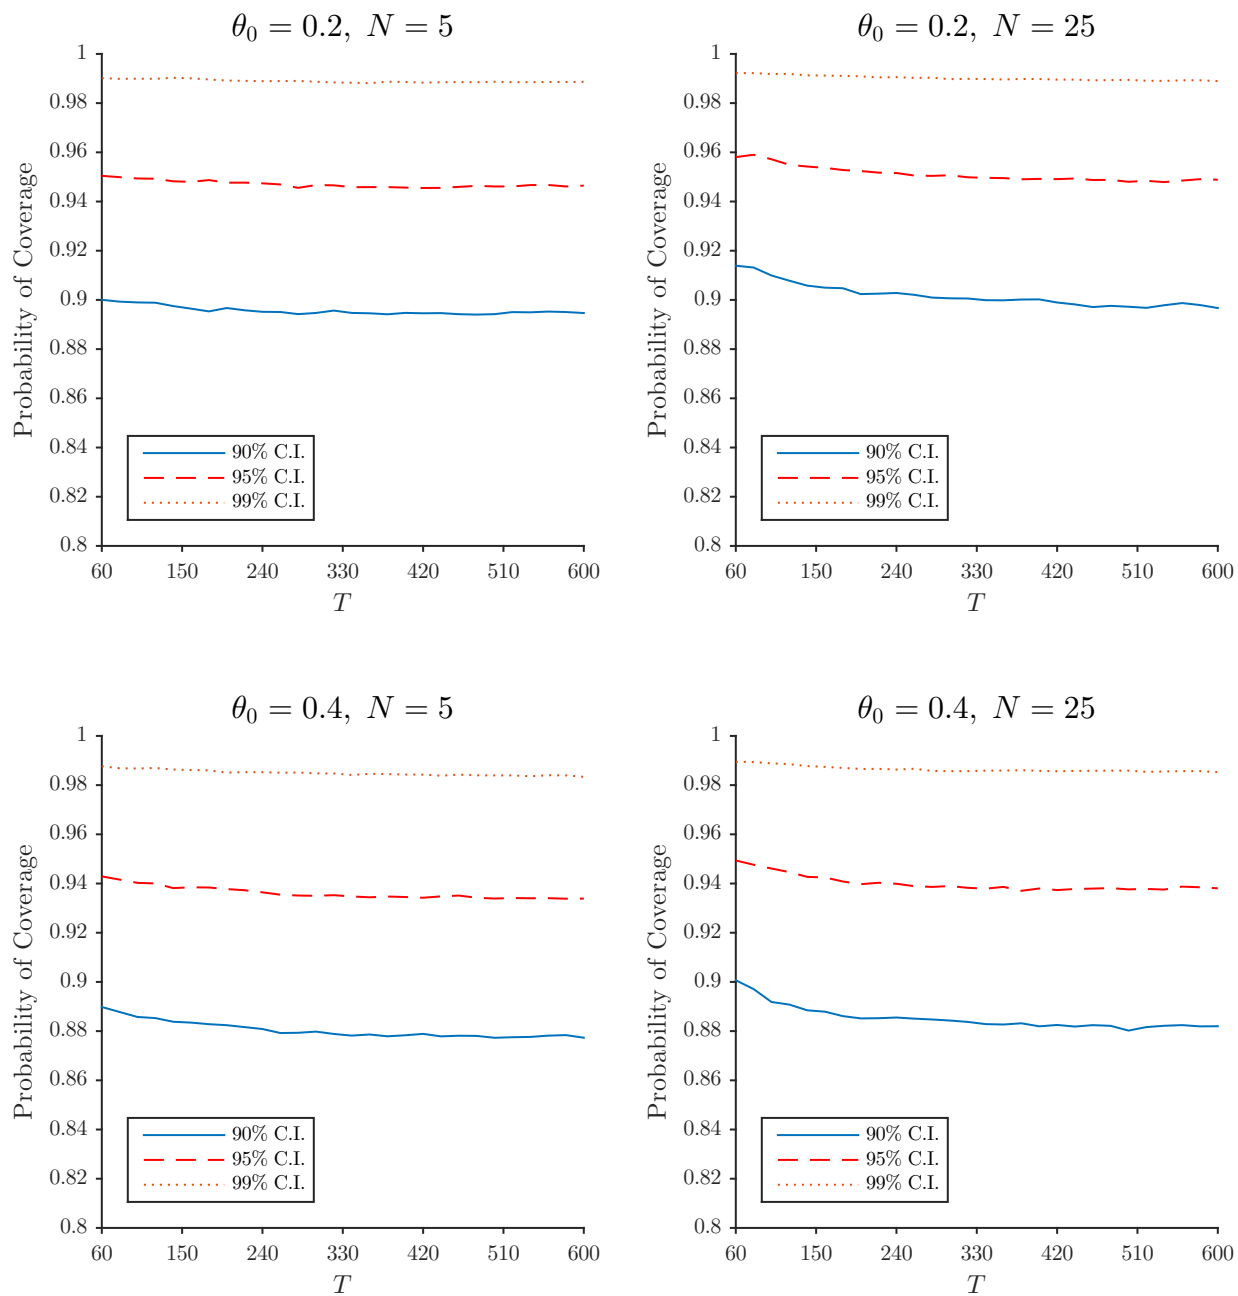


Figure 13 Probabilities of Coverage of the Confidence Intervals for the Constrained Hansen-Jagannathan Bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the constrained HJ-bound for different number of risky assets (N) and length of time series 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.

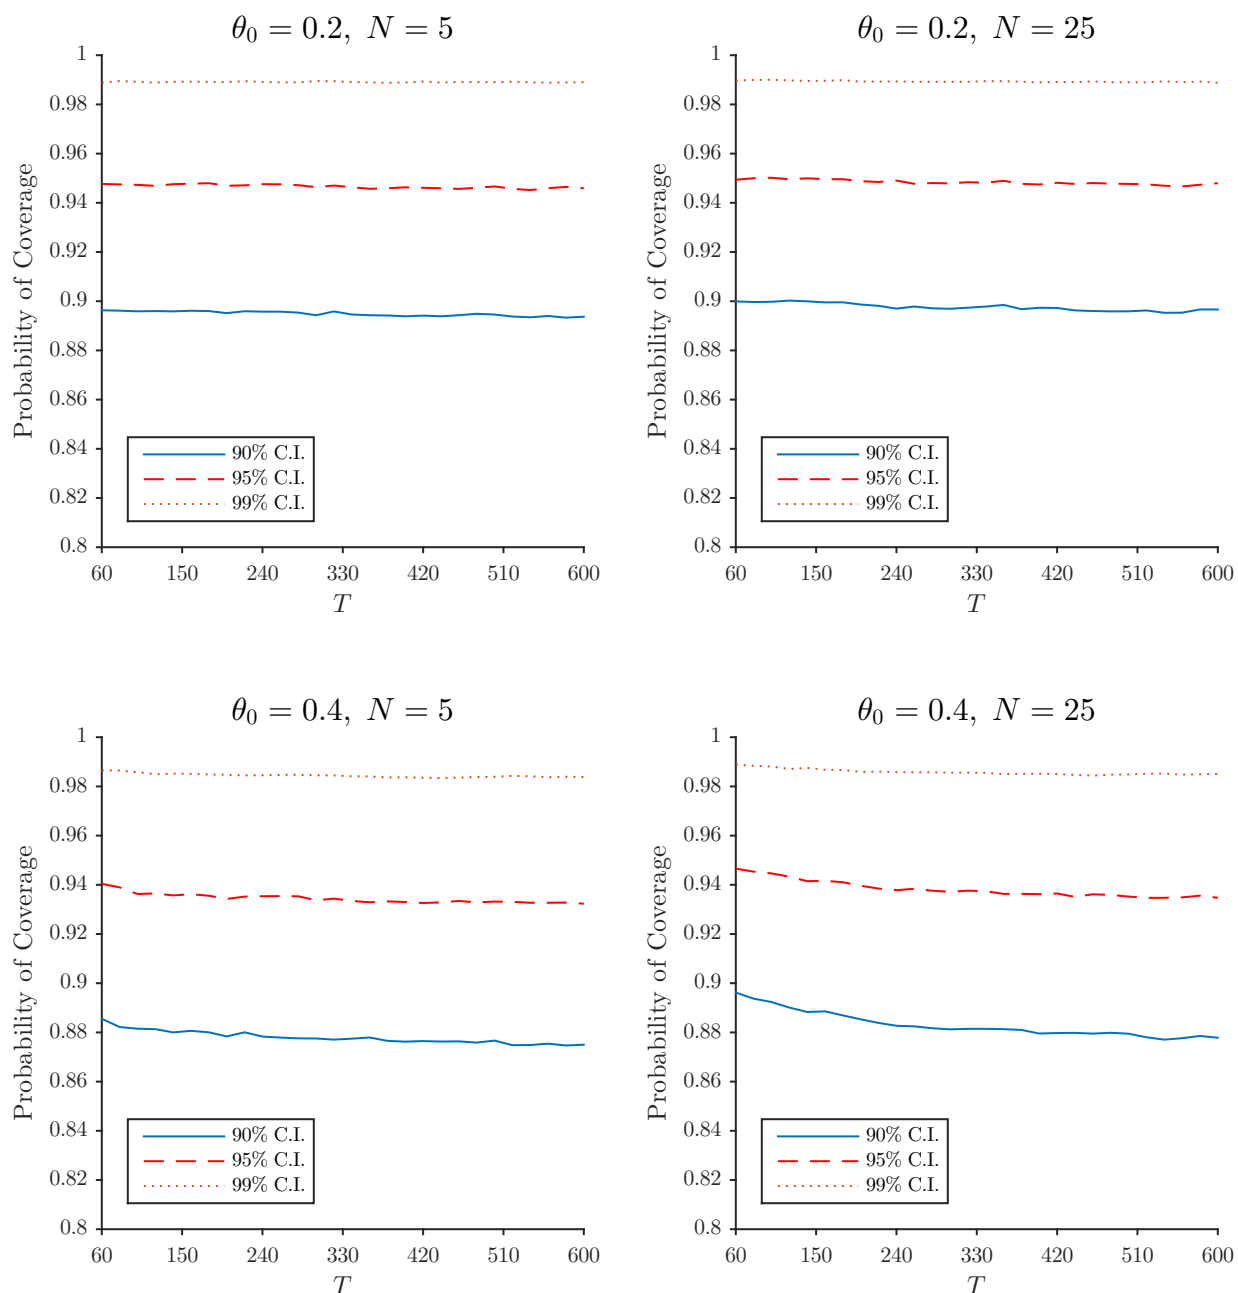


Figure 14 Probabilities of Coverage of the Confidence Intervals for the Unconstrained Hansen-Jagannathan Bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the unconstrained HJ-bound for different number of risky assets (N) and length of time series 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 1946/1–2014/10. 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.

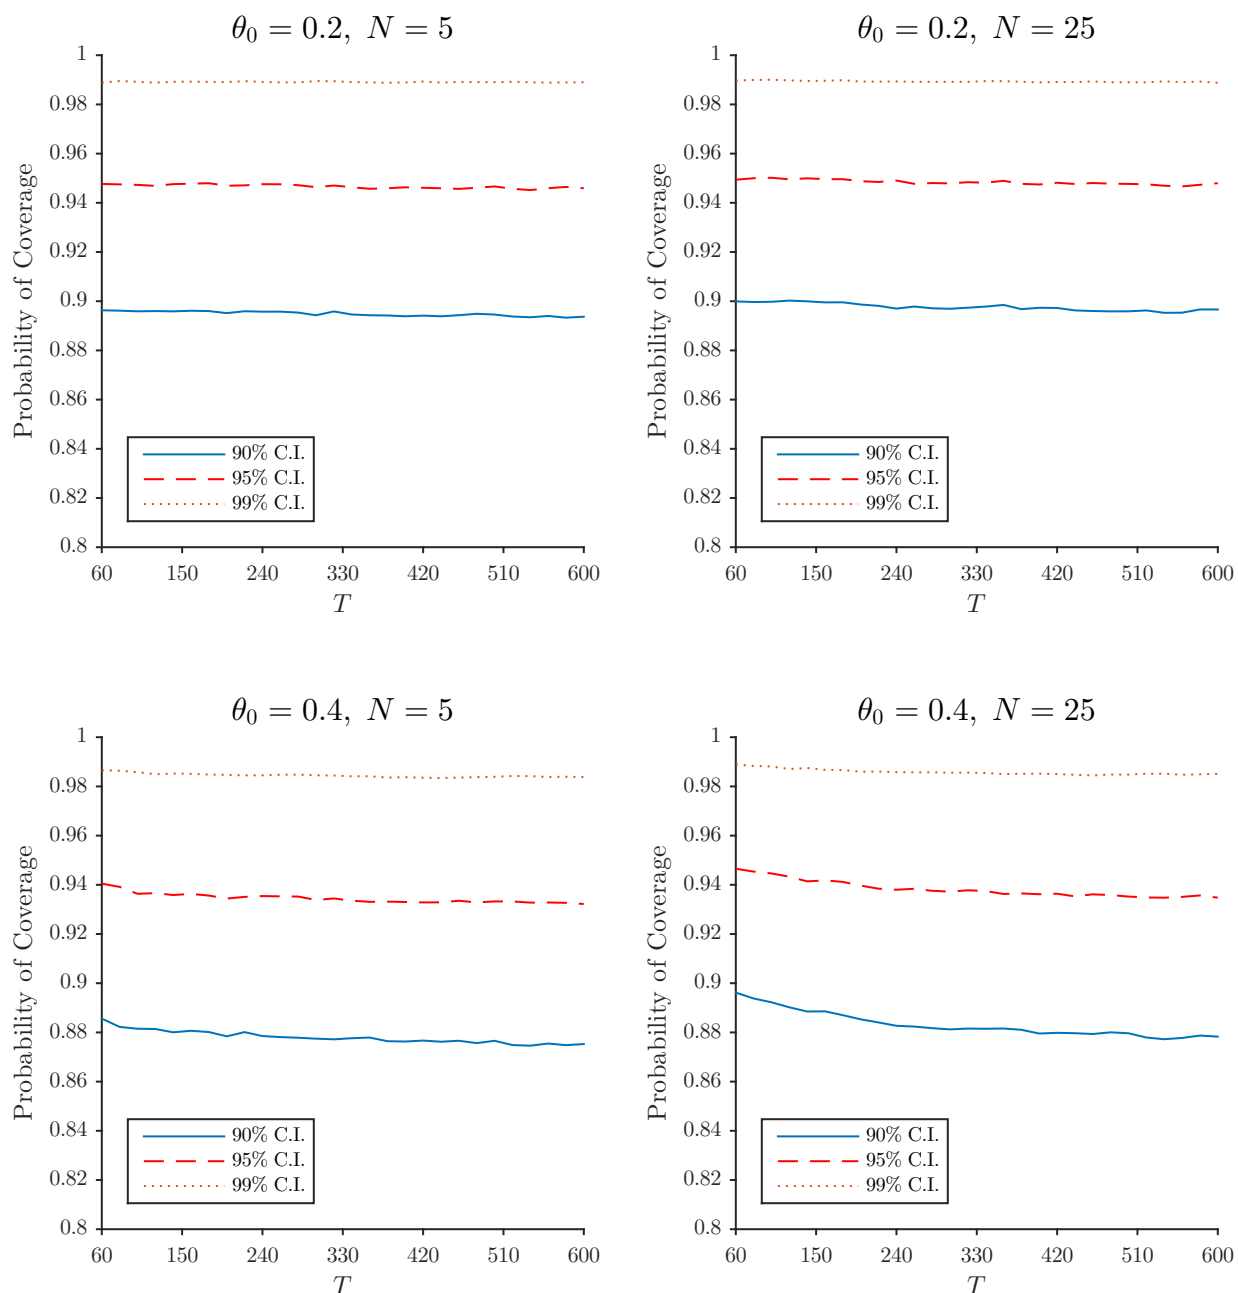


Figure 15 Probabilities of Coverage of the Confidence Intervals for the Constrained Hansen-Jagannathan Bound. The figure presents the probabilities of coverage of the 90%, 95%, and 99% confidence intervals for the constrained HJ-bound for different number of risky assets (N) and length of time series 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 1946/1–2014/10. 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 16 Probability Density Functions of the Sample Hansen-Jagannathan Bounds.

The figure presents the estimated pdfs of the sample unconstrained HJ-bound and of the nonparametric estimator of the constrained HJ-bound. The parameters of the distributions are estimated based on the monthly returns on the 25 Fama-French size and book-to-market ranked portfolios over the period 1931/7–2014/10. When not available analytically, the pdfs based on the finite-sample distributions are estimated using 100,000 simulations.