The Opportunity Set: *Market Opportunities and the Effective Breadth of a Portfolio*

RICHARD C. GRINOLD AND MARK P. TAYLOR

RICHARD C. GRINOLD
is global head of research at Barclays Global Investors in San Francisco, CA.
richard.grinold@barclaysglobal.com

MARK P. TAYLOR
is a managing director of Barclays Global Investors, London, and professor of international finance at the University of Warwick in Coventry, U.K.
mark.taylor@barclaysglobal.com

The opportunity set (OS) is an ex post measure of the maximum Sharpe ratio that could have been obtained in a market for any portfolio constructed on a given universe of assets.

In this article, we show that the average value of the OS is approximately equal to the square root of portfolio breadth (number of assets). We, therefore, suggest an interpretation of the OS as the *effective breadth* of a portfolio. Effective breadth is higher when the volatility of asset returns is abnormally high. In such an environment, the market will offer abnormal opportunities for a long/short investor, which is tantamount to having more portfolio breadth. Conversely, when asset returns volatility is extremely low, even the most skillful long/short investor would be challenged to generate excess returns for any given level of risk. In such an environment, it is as if there is little breadth at all in the portfolio. The effective breadth interpretation continues to hold even if we view the returns as residual to some reference portfolio that captures the general movement of the universe of assets.

In two appendices, we provide further details and extensions of our opportunity set analysis and also derive the statistical distribution of the opportunity set through a set of Monte Carlo experiments under a range of assumptions concerning the statistical distribution of returns.

THE OPPORTUNITY SET

The opportunity set is defined as $OS \equiv \sqrt{r'\Omega^{-1}r}$, where r is the vector of ex post excess returns to the assets of a portfolio and Ω is the covariance matrix of those returns (Grinold [2006]). The opportunity set earns its name from the fact that OS measures the opportunities that were available in the market to long/short investors. Suppose, for example, that the portfolio had only one asset (e.g., the market) with return variance ω^2, so that $OS \equiv \sqrt{r^2/\omega^2}$. Now suppose, on the one hand, that $r^2 > \omega^2$, so that $OS > 1$; that is, actual returns volatility was greater than the expected volatility of returns, meaning that exceptional opportunities were available to an investor who was either long the market and short cash, or the converse. If, on the other hand, the OS was very small, meaning that there was little movement in market returns, even a correctly positioned investor could expect little profit for an acceptable degree of ex ante risk. On average, however, we would expect the average returns volatility to be equal to the variance of returns, so that $OS \approx 1$.

The Opportunity Set as Effective Breadth

As Appendix A shows, this result generalizes to the N-asset case. With N assets in our portfolio, we expect, on average, that the size of the opportunity set will be approximately equal to the square root of the number of assets in the portfolio:

$$E(OS) \approx \sqrt{N} \tag{1}$$

Thus, if the realized OS is greater than $\sqrt{N}$, the market offered exceptionally good opportunities to a long/short investor, but if the OS was very much less than $\sqrt{N}$, the market offered relatively poor opportunities for adding client value.[1]

Because the average value of the opportunity set is related to the number of assets, or breadth of the portfolio, it is informative to think of OS as the effective breadth of the portfolio. When a large amount of variation in market returns exists, a skillful long/short investor can add significant client value in a way that is similar to expanding the portfolio by adding independent assets to it. Suppose, however, that there was no movement at all in returns (i.e., $r = 0$); then it as if there were no breadth at all (i.e., no assets) because no amount of skill can profit in a static market.

If r_p is the ex post portfolio excess return on portfolio P, then it can also be shown (see Appendix A) that

$$r_p = OS \cdot IC_p \cdot \omega_p \tag{2}$$

where OS is the size of the opportunity set, IC_p is the realized information coefficient of portfolio P, and ω_p is the predicted level of risk in the portfolio. This ex post result is analogous to the ex ante result (for residual returns),

$$\alpha_p = IR \cdot TC_p \cdot \omega_p \tag{3}$$

where IR is the information ratio, $IR = \sqrt{\alpha'\Omega^{-1}\alpha}$ (see Grinold and Kahn [2000]); TC_p is portfolio P's transfer coefficient, as in Clarke, DeSilva, and Thorley [2002]; and ω_p is the predicted level of residual risk in the portfolio.

Because IC_p is a correlation, it is less than or equal to one; thus, for any portfolio P, the opportunity set gives an upper bound on the amount of return we can obtain per unit of risk taken:

$$r_P / \omega_P \leq OS \tag{4}$$

As shown in Appendix A, we can always find a portfolio (*hindsight portfolio*) that will achieve the bound defined in Equation (4). Another immediate—if somewhat straightforward—insight offered by the return decomposition expressed in Equation (2) is that while a small opportunity set and small risk budget may be at least partly responsible for a low portfolio return, negative portfolio return can only be due to a shortfall in skill. This follows because, by definition, we have $OS \geq 0$, $1 \geq IC_p \geq -1$, and $\omega_p \geq 0$. Thus, $r_p < 0$ if, and only if, $IC_p < 0$.

The result $E(OS) \approx \sqrt{N}$ does not require that asset returns be uncorrelated. Neither does the result hinge on any distributional assumptions concerning returns. This is, therefore, a strong result.

It is possible to split the opportunity set into a portion due to the motion of the market as a whole and a portion due to the residual returns. We can use any reference portfolio as a guide. In the examples that follow, we use a capitalization-weighted universe of assets in one case and equal weights in another.

Call the reference portfolio U with holdings $u = \{u_1, u_2, \ldots, u_N\}'$. This portfolio has variance $\omega_U^2 = u'\Omega u$. The return on the reference portfolio is $r_U = r'u$. If we define beta as $\beta = \Omega u / \omega_U^2$, then we can split the returns into market-wide and residual components, as follows:

$$r = \beta r_U + \theta \tag{5}$$

This allows us to also split the opportunity set into market-wide and residual components,

$$r'\Omega^{-1}r = \frac{r_U^2}{\omega_U^2} + \theta'\Omega^{-1}\theta \tag{6}$$

Use the notations $ROS = \sqrt{\theta' \cdot \Omega^{-1}\theta}$ and $MOS = |r_U| / \omega_U$ to distinguish the residual opportunity set and market opportunity set, respectively, from $OS = \sqrt{r' \cdot \Omega^{-1} r}$. With this terminology, Equation (6) becomes

$$OS^2 = MOS^2 + ROS^2 \tag{7}$$

As we show in Appendix A, the expected value of OS^2 is approximately equal to N, the number of assets. The expected value of MOS^2 is close to one, so that the expected size of the residual opportunity set squared, ROS^2, must be close to $N - 1$.

The analogue of Equation (4) is also valid for the residual opportunity set. For any portfolio P with residual return θ_p, we have the bound

$$\theta_p / \omega_p \leq ROS \tag{8}$$

In addition, it turns out that the residual holdings of the hindsight portfolio will achieve that bound.[2]

The following examples show that the ROS and OS are practically indistinguishable in a market with a large number of assets.

Examples. Our first example is from the U.S. large-capitalization market for the years 1985 through 2007. The data are monthly for 276 months. The BARRA risk model was used to compute the opportunity set. The number of assets averaged 1,103 during the period and varied from a low of 999 in April 1990 to a high of 1,176 in November 1996. Extremely high returns were screened out.[3]

Exhibit 1 shows the histogram of $OS(t)/\sqrt{N(t)}$ for $t = 1, \ldots, 276$.

The average of these realizations is 0.96—as expected, slightly less than one—and the standard deviation is 0.11. The maximum OS occurred in September 1998 and the minimum in December 2002. The MOS peaked at 4.93 in the market-crash month of October 1987.[4] The $OS/\sqrt{N}$ and $ROS/\sqrt{N-1}$ are very close to one for that crash month, indicating there was no special opportunity to separate winners and losers; they were all losers. In Exhibit 2 we present some summary statistics.

The first conclusion we can draw from Exhibit 2 is that OS and ROS are practically indistinguishable when there are a large number of assets. Indeed, the correlation between the two series is 0.99997. The second conclusion is that the average value of $OS/\sqrt{N}$ is slightly less than one, as we should expect. The third conclusion is that skewness and (excess) kurtosis of the MOS series are not consistent with normality; indeed, the MOS series fails the normality test in quite spectacular fashion.[5] A fourth point is the significant degree of autocorrelation at one- and two-month lags in the OS series; that is, a larger (smaller) than expected opportunity set in one month will most likely be followed by a larger (smaller) opportunity set in the next month. The opportunity set was 30% larger in September 1998 than would be expected on average, compared to the opportunity set in December 2002, which was only 67% of its expected level.

Our second example comes from investing in the aggregate stock markets of 13 developed countries: Australia, Japan, Hong Kong, U.S., Canada, U.K., France, Germany, Italy, Spain, Sweden, Switzerland, and the Netherlands. The data are monthly from January 1976 through June of 2006 (367 months). We used an in-sample covariance matrix for the Ω used to calculate the opportunity set. Exhibit 3 contains the histogram of the residual opportunity set scaled by the square root of 12 (the number of assets minus 1).

EXHIBIT 1

Scaled Opportunity Set Histogram, U.S. Large-Capitalization Market, 1985–2007

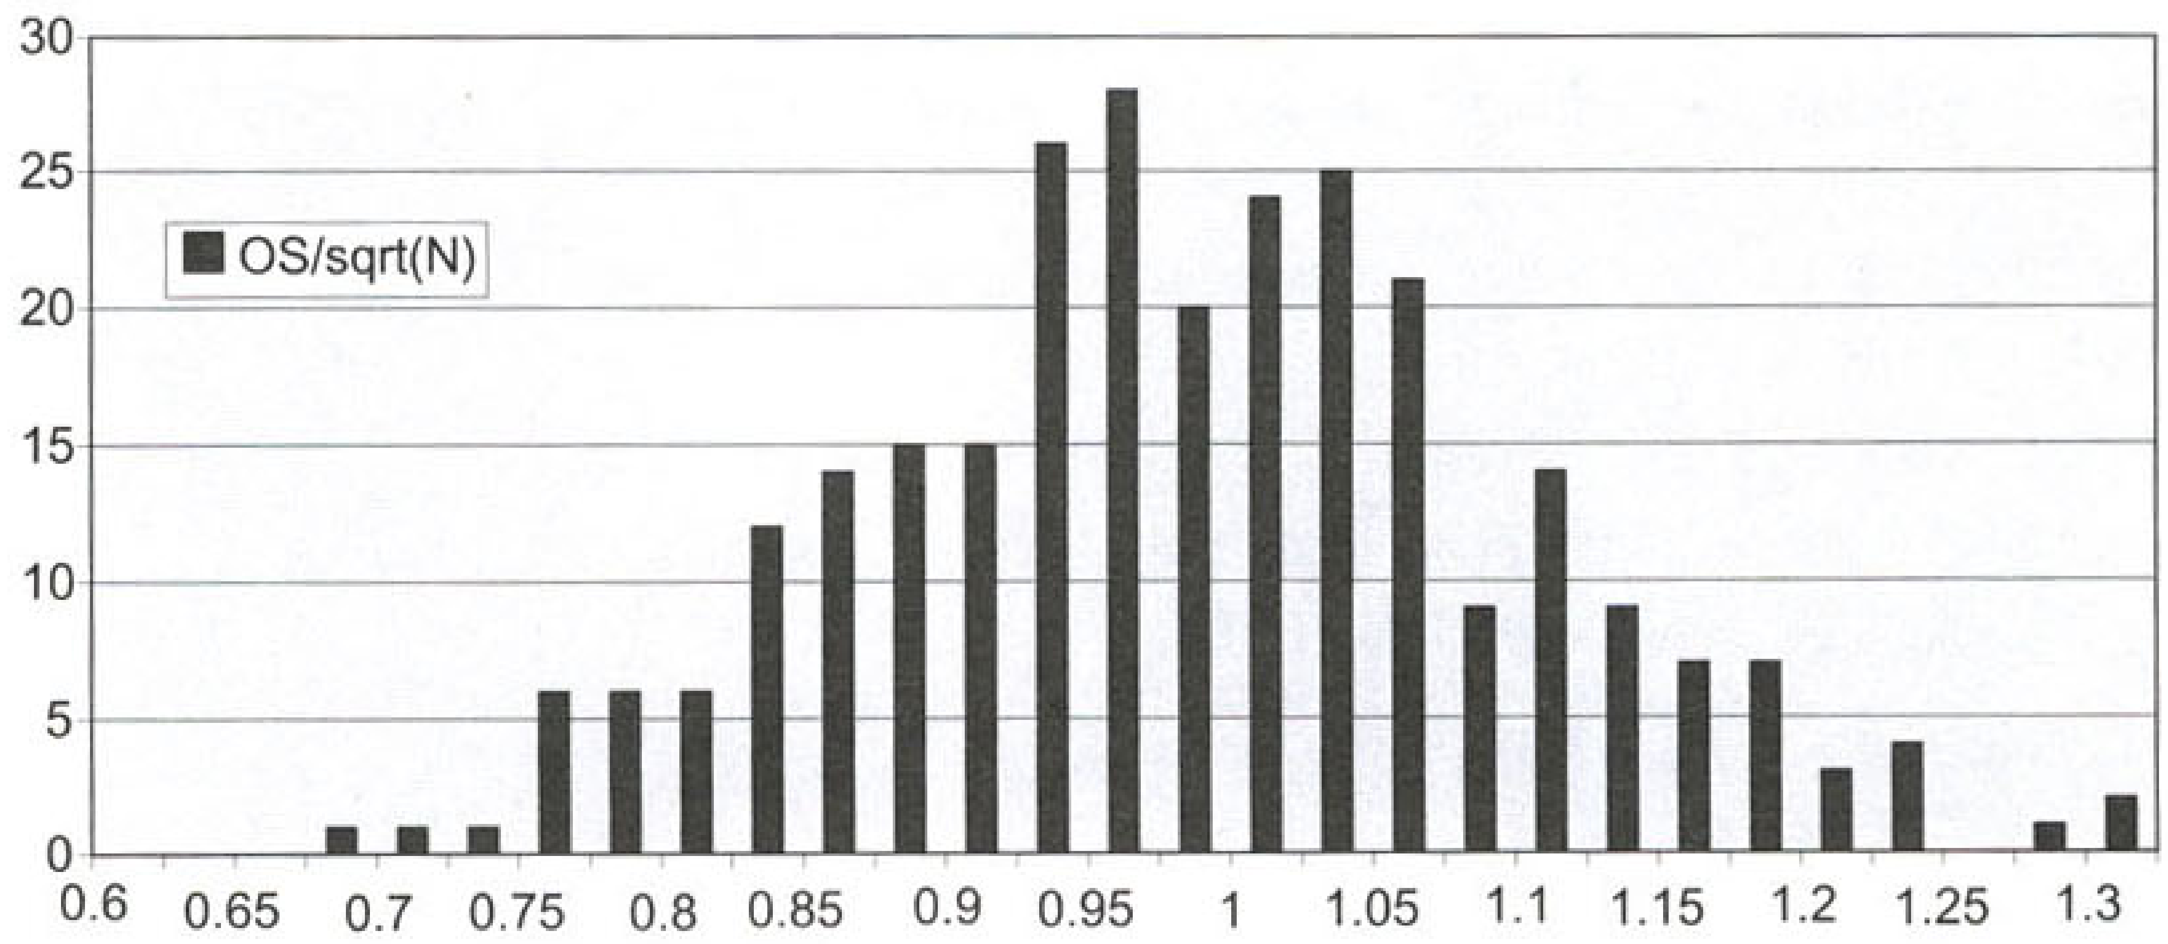

EXHIBIT 2

Opportunity Set Summary Statistics, U.S. Large-Capitalization Market, 1985–2007

	MOS	$OS/\sqrt{N}$	$ROS/\sqrt{N-1}$
Average	0.809	0.985	0.984
Standard deviation	0.630	0.125	0.125
Maximum	0.4928	1.608	1.607
Minimum	0.003	0.673	0.673
1st order autocorrelation	–0.023	0.279	0.278
2nd order autocorrelation	0.014	0.142	0.141
Skew	2.030	0.487	0.491
Kurtosis	8.197	1.672	1.686

The distribution is quite wide. In 20% of the months, the *ROS* is less than 65% of its expected value, and at the other end of the distribution, in 20% of the months, the *ROS* is 125% or more of the expected value.

The summary statistics for the *OS* and *ROS* for this example are displayed in Exhibit 4. The *MOS* and $OS/\sqrt{N}$ reach their maxima (6.16 and 2.43, respectively) in the stock market-crash month of October 1987. The $ROS/\sqrt{N-1}$ hits a respectable 1.79 in the crash month but does even better in five other months with a peak value of 2.09 in January 1987.

As in the previous example, there is considerable autocorrelation in the *OS* and *ROS*. The scaled averages are, as anticipated, slightly less than one. The correlation between the *OS* and *ROS* series in this sample is 0.98. There is again significant departure from normality in each of the series. We examine the implications of departures from normality in Appendix B.

The Opportunity Set and Risk Predictions

The *OS* formula has two inputs: the returns, *r*, and the risk prediction Ω. The question can always be asked: Is the *OS* large because the returns are large or because the risk prediction was too low? In any single period, we cannot give a definitive answer. One way to address the question is to look for a pattern of consecutive high or low values for the *OS*. Evidence of such a pattern would be indicative that the predicted covariance matrix is not up to the job.

We can see this phenomenon quite clearly in the second of the examples described in the previous section. As noted, we used the full sample covariance matrix to construct the *OS*, which has the advantage of being correct on average. The risk of the markets, however, was considerably higher than average in the earlier part of the sample history and considerably lower than average over the more recent history. One way to see this is to cumulate the values of $OS(t)/\sqrt{N}$ less a benchmark value of one per period. The series is

EXHIBIT 3

Scaled Opportunity Set Histogram, 13 Aggregate Markets, 1976–2006

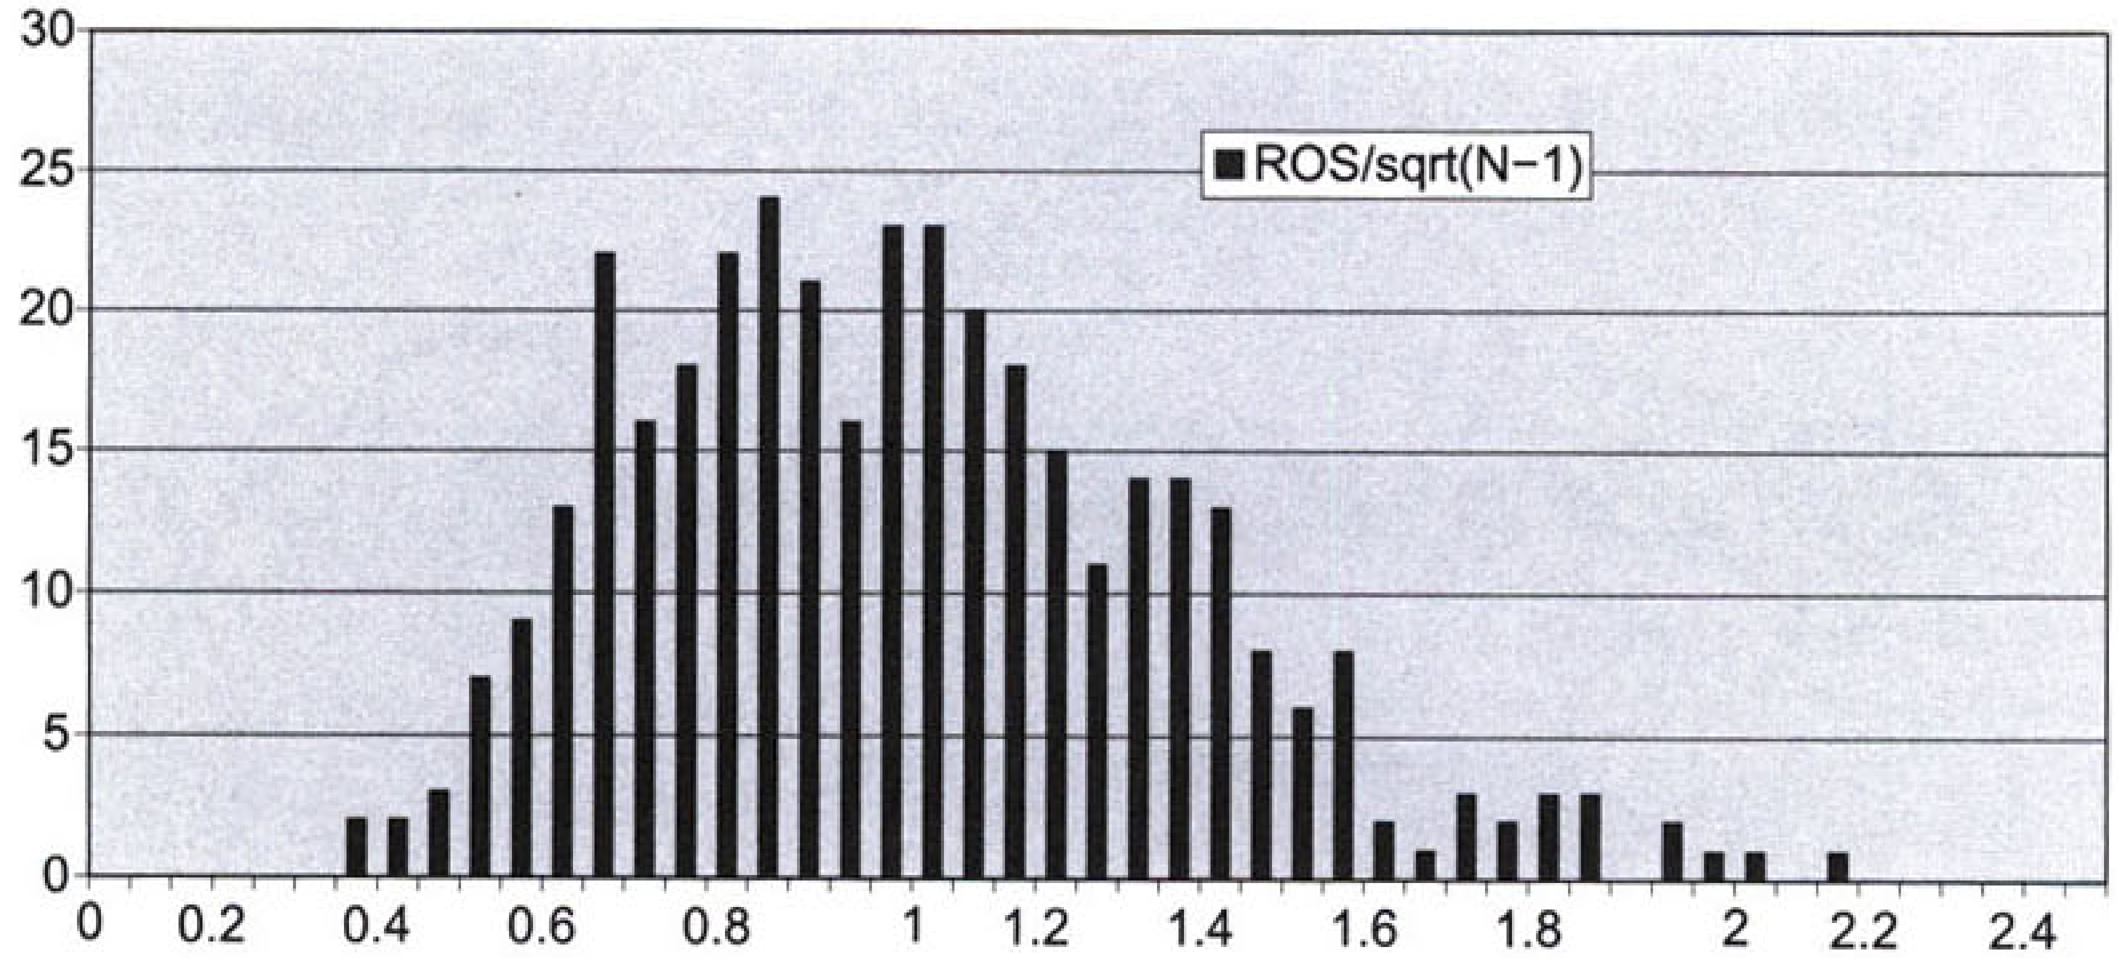

EXHIBIT 4
Opportunity Set Summary Statistics, 13 Aggregate Markets, 1976–2006

	MOS	$OS/\sqrt{N}$	$ROS/\sqrt{N-1}$
Average	0.740	0.950	0.940
Standard deviation	0.681	0.330	0.330
Maximum	6.161	2.430	2.090
Minimum	0.000	0.260	0.270
1st order autocorrelation	0.110	0.460	0.500
2nd order autocorrelation	0.060	0.371	0.411
Skew	2.400	0.760	0.600
Kurtosis	11.600	0.980	0.220

$$X(t) = \left\{ \frac{1}{\sqrt{N}} \cdot \sum_{\tau=1,T} OS(\tau) \right\} - t \tag{9}$$

Exhibit 5 provides the chart of that series for our second example.

The low-risk predictions persist until mid-1988. From that point on, the risk prediction is too high. The end value is consistent with the average 0.95 over the 367 months. This observation is included to point out that another good reason to monitor the opportunity set is as a monitor of risk predictions.

CONCLUSION

The opportunity set, defined as $OS \equiv \sqrt{r'\Omega^{-1}r}$ measures, is in an ex post sense the opportunities that were available in the market to investors. If, for example, ex post returns volatility was high relative to the covariance matrix of returns, then a long/short investor was offered exceptional opportunities. Whether the active manager was able to capitalize on this depends on the level of investment skill and the aggressiveness of the portfolio.

In this article, we have shown that $E(OS^2) = N$, where N is the number of assets in the portfolio. This result holds, moreover, regardless of the distribution of asset returns. Further, this result implies that $E(OS) \approx \sqrt{N}$. The intuitive interpretation of this result is that the opportunity set measures the effective breadth of a portfolio at a given point in time. So when the market offered abnormally high opportunities in terms of the volatility of realized asset returns, it was as if the breadth of the portfolio were expanded ($OS = \sqrt{r'\Omega^{-1}r} > \sqrt{N}$). Conversely, when there was shrinkage in returns volatility in a given period, it was as if the breadth of the portfolio had shrunk ($OS = \sqrt{r'\Omega^{-1}r} < \sqrt{N}$).

In Appendix A, we derive and describe the properties of the opportunity set in more detail. In Appendix B, we describe a set of Monte Carlo experiments in which residual returns are generated by a variety

EXHIBIT 5
Cumulative Scaled and De-Meaned Opportunity Set, 13 Aggregate Markets, 1976–2006

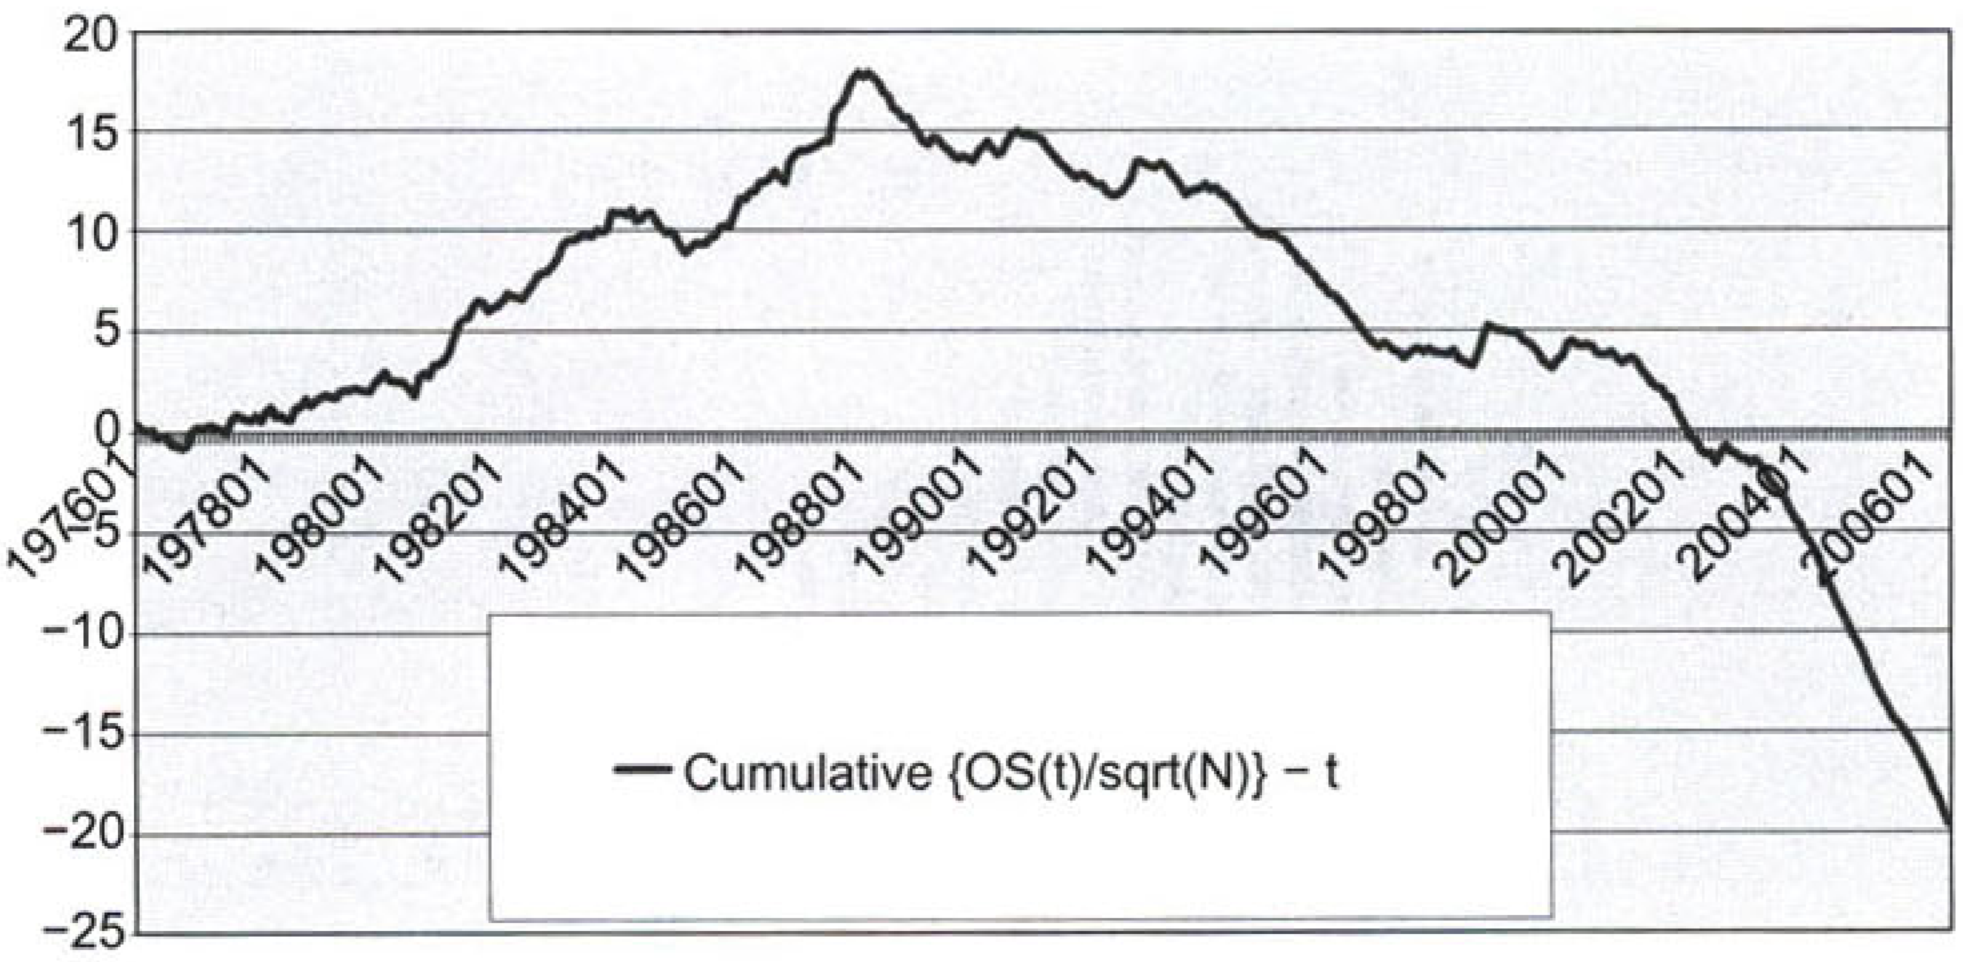

of multivariate distributions, ranging from multivariate normal to highly leptokurtic multivariate Student's t. The experiments illustrate the result that $E(OS^2) = N$ is distribution independent. They also show that, on the one hand, while the approximation $E(OS) \approx \sqrt{N}$ continues to hold for a range of degrees of leptokurtosis, the inequality $E(OS) < \sqrt{N}$ becomes more marked for greater degrees of kurtosis. On the other hand, the variance of OS is shown to be higher for more fat-tailed return distributions and approaches a lower bound of 0.5 as the distribution approaches normality.

Appendix A

The Opportunity Set and Effective Portfolio Breadth

Consider a universe of N assets and an $N \times N$ covariance matrix Ω whose (n, m)th element is the predicted, or ex ante, covariance of the return to the nth and mth assets. Consider any portfolio P with holdings $p = \{p_1, p_2, \ldots, p_N\}'$. The risk of this portfolio is

$$\omega_p = \sqrt{p'\Omega p} \tag{A1}$$

Now consider the returns on the assets $r = \{r_1, r_2, \ldots, r_N\}'$. We can associate a portfolio S with holding $s = \{s_1, s_2, \ldots, s_N\}'$ with these returns through the equation,

$$r = \Omega s \tag{A2}$$

The holdings in S are proportional to the holdings the investor wished he had held (i.e., the perfect *hindsight portfolio*).

Notice that the risk of portfolio S is equal to the size of the opportunity set,

$$s'\Omega s = r'\Omega^{-1} r \tag{A3}$$

The covariance between portfolios P and S is

$$p'\Omega s = p'r = \sqrt{s'\Omega s} \cdot \left\{ \frac{p'\Omega s}{\sqrt{s'\Omega s} \cdot \sqrt{p'\Omega p}} \right\} \cdot \sqrt{p'\Omega p} \tag{A4}$$

From Equations (A1) and (A3), Equation (A4) becomes

$$p'r = \sqrt{r'\Omega^{-1} r} \cdot \left\{ \frac{p'\Omega s}{\sqrt{s'\Omega s} \cdot \sqrt{p'\Omega p}} \right\} \cdot \omega_p$$

or

$$r_p = \sqrt{r'\Omega^{-1} r} \cdot IC_p \cdot \omega_p \tag{A5}$$

The first multiplicative term on the right side of Equation (A5) is the opportunity set, first defined by Grinold [2006],

$$OS \equiv \sqrt{r'\Omega^{-1} r} \tag{A6}$$

The second term on the right side of Equation (A5) is the correlation between the portfolios S (the positions the investor wished he had held) and P (the positions the investor actually held). We call this correlation the *realized information coefficient* of portfolio P, or IC_p. Because IC_p is a correlation coefficient, it is weakly less than one. Therefore, for any portfolio P,

$$r_P / \omega_P \leq \sqrt{r'\Omega^{-1} r} = r_S / \omega_S \tag{A7}$$

The opportunity set is thus a tight upper bound on the return per unit of risk that can be obtained in any portfolio.

The realized IC is akin to the transfer coefficient that is an ex ante measure of how close the positions held in portfolio P are to those the investor would like to have held, but are prevented from holding due to constraints, transactions costs, and the like.

Suppose that $\alpha = \{\alpha_1, \alpha_2, \ldots, \alpha_N\}'$ is our forecast of alpha. We can link alpha to a portfolio Q with holdings $q = \{q_1, q_2, \ldots, q_N\}'$ in the spirit of Equations (A2) and (A3)

$$\begin{aligned} \alpha &= \Omega q \\ q'\Omega q &= \alpha'\Omega^{-1}\alpha = IR^2 \end{aligned} \tag{A8}$$

Then the analogue to Equation (A5) is

$$\begin{aligned} \alpha_p = p'\alpha &= \sqrt{\alpha'\Omega^{-1}\alpha} \cdot \left\{ \frac{p'\Omega q}{\sqrt{q'\Omega q} \cdot \sqrt{p'\Omega p}} \right\} \cdot \omega_p \\ &= IR \cdot TC_p \cdot \omega_p \end{aligned} \tag{A9}$$

The correlation in curly brackets in Equation (A9) is the transfer coefficient.

Note that OS is a pure number (i.e., dimensionless), which is clear from its definition, $OS \equiv \sqrt{r'\Omega^{-1}r}$, but also follows from the decomposition $r_p = OS \cdot IC_p \cdot \omega_p$. Because r_p and ω_p are both measured in percentage terms, and IC_p is a correlation coefficient and hence a pure number, by matching the dimensions on either side of the definition, it follows that OS must also be a pure number.[6]

Further insight into the nature of the opportunity set may be gained by noting that $E(OS) = E\left[\sqrt{r'\Omega^{-1}r}\right] \approx \sqrt{E(r'\Omega^{-1}r} = \sqrt{N}$, where N is the number of assets in the portfolio, and hence is the dimension of Ω and r. Consider, first, the case in which the expected return on the assets is zero, then $E(rr') = \Omega$. In this case we can establish that $E(r'\Omega^{-1}r) = N$. To see this, note that $E(r'\Omega^{-1}r)$ is a scalar,[7]

$$\begin{aligned} E(OS^2) &= E(r'\Omega^{-1}r) = trace[E(r'\Omega^{-1}r)] \\ &= trace[E(\Omega^{-1}rr')] = trace[\Omega^{-1}E(rr')] \\ &= trace[\Omega^{-1}\Omega] = trace[I_N] = N \end{aligned} \quad \text{(A10)}$$

where I_N is the N-dimensional identity matrix.

In the more general case, we have

$$E(rr') = \Omega + \mu\mu' \quad \text{(A11)}$$

where μ is the expected excess return. We will argue that the trace of $\Omega^{-1}E(rr')$ is still approximately equal to N.[8]

To see this, define a portfolio M with holdings $m = \{m_1, m_2, \ldots, m_N\}$, risk $\omega_M = \sqrt{m'\Omega m}$, and expected excess return, $\mu_M = \mu' m$, such that

$$\mu = \left\{\frac{\Omega m}{m'\Omega m}\right\} \bullet \mu_M \quad \text{(A12)}$$

M is the portfolio with the highest Sharpe ratio, $SR_M = \mu_M/\omega_M$. Let $\beta \equiv \Omega\, m/\omega_M^2$. Then,

$$\mu = \beta \cdot \mu_M = \left\{\frac{\Omega m}{m'\Omega m}\right\} \cdot \mu_M = \Omega m \cdot SR_M \cdot \frac{1}{\omega_M} \quad \text{(A13)}$$

and

$$\mu\mu' = \Omega \cdot (m\beta') \cdot SR_M^2 \quad \text{(A14)}$$

Thus, from Equations (A10) and (A13),

$$\Omega^{-1}E(rr') = \Omega^{-1}\left\{\Omega + \mu\mu'\right\} = I_N + SR_M^2 \cdot (m\beta') \quad \text{(A15)}$$

The trace is a linear operator, so

$$\begin{aligned} trace\{I_N + SR_M^2 \cdot (m\beta')\} &= trace(I_N) + SR_M^2 \cdot trace(m\beta') \\ &= N + SR_M^2 \cdot trace(m\beta') \\ &= N + SR_M^2 \end{aligned} \quad \text{(A16)}$$

Hence, from Equations (A14) and (A15),

$$E(OS^2) = E(r'\Omega^{-1}r) = trace[\Omega^{-1}E(rr')] = N + SR_M^2 \quad \text{(A17)}$$

so that

$$E(OS^2)/N = 1 + SR_M^2/N \quad \text{(A18)}$$

Thus, the closeness of the approximation $E(OS) \approx \sqrt{N}$ is seen to be a decreasing function of SR_M^2/N.[9] Suppose, however, we take a low value of N equal to 10 and a high value of the expected Sharpe ratio equal to 0.5. We would then have $E(OS^2) = 10.25 \approx 10$ and the approximation will still be close. For the range of values of N and SR_M^2 more usually encountered, the approximation would be even closer.

Overall, therefore, we have demonstrated in this appendix that the unconditional expected value of the opportunity set is approximately equal to the breadth of the portfolio in terms of the number of assets. The intuitive interpretation of this result is that the opportunity set in some sense measures effective breadth. It will be close to the square root of the number of assets N on average, but may be higher or lower in any particular sample. For example, if rr' is large relative to the covariance matrix Ω, such that $OS \equiv \sqrt{r'\Omega^{-1}r} > \sqrt{N}$, then the market offered abnormal opportunities in the period for which r is calculated, which is tantamount to having more breadth. At the other extreme, suppose the market return vector was null, then $OS = 0$, and it is as if there were no breadth at all in the portfolio (i.e., no assets).

Appendix B

Jensen's Inequality and the Distribution of the Opportunity Set

Appendix A shows that $E(OS^2) = N$ exactly when the expected excess return is zero. The approximation $E(OS) \approx \sqrt{N}$ follows from this relation because of Jensen's inequality: if

X is a real-valued random variable with $E(|X|)$ finite and the function $g(X)$ is concave, then $E[g(X)] \le g(E[X])$. Thus, $E(OS) = E\left[\sqrt{r'\Omega^{-1}r}\right] \le \sqrt{E(r'\Omega^{-1}r)} = \sqrt{N}$. To get a feel for the empirical importance of this effect, we carried out a set of Monte Carlo experiments. We initially assumed that returns were normally distributed and then relaxed this assumption.

Assuming Normally Distributed Returns

First, note that if we are willing to assume that returns are normally distributed, then $OS^2 = r'\Omega^{-1}r$ is a quadratic form in an N-dimensional vector of normally distributed random variables and is therefore distributed as χ^2 with N degrees of freedom. Because the mean of a χ^2 distribution is equal to its degrees of freedom (in this case N), we have $E(OS^2) = N$, which again implies $E(OS) \approx \sqrt{N}$ (although our earlier derivation of this result was distribution free).

In order to keep our Monte Carlo experiments manageable, we assumed a relatively small portfolio size in terms of the number of distinct assets. One area of active management that naturally has a relatively small number of assets is active currency management. It is thus possible to think of a currency portfolio in which there might be, say, 10 currencies and 9 effective assets after 1 currency is chosen as a numeraire.[10] In order to calibrate the data-generating mechanism for the Monte Carlo experiments, we estimated a 9 × 9 covariance matrix, Ω, using daily foreign exchange returns expressed in U.S. dollars, for the Australian dollar, Canadian dollar, Euro, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, Swiss franc, and pound sterling over the period January 1999 to March 2007. We then drew r from a multivariate normal distribution with covariance matrix Ω and calculated $OS = \sqrt{(r'\Omega^{-1}r)}$, repeating this for a total of 100,000 draws.

Summary statistics for the resulting empirical distribution of OS are shown in Exhibit B1. As anticipated, $E(OS)$ is indeed less than $\sqrt{N} = \sqrt{9} = 3$, but only slightly so with $E(OS) = 2.92 < 3$.

The resulting empirical distribution of the opportunity set, $OS \equiv \sqrt{(r'\Omega^{-1}r)}$, is drawn in Exhibit B2.

EXHIBIT B1

Statistics for the Empirical Distribution of $OS = \sqrt{(r'\Omega^{-1}r)}$ for $r \sim N(0, \Omega)$, $N = 9$

E(OS)	Median	5th Percentile	25th Percentile	75th Percentile	95th Percentile	Variance
2.92	2.89	1.83	2.43	3.38	4.10	0.48

EXHIBIT B2

Empirical Distribution of the Opportunity Set, $OS = \sqrt{(r'\Omega^{-1}r)}$, for $r \sim N(0, \Omega)$, $N = 9$

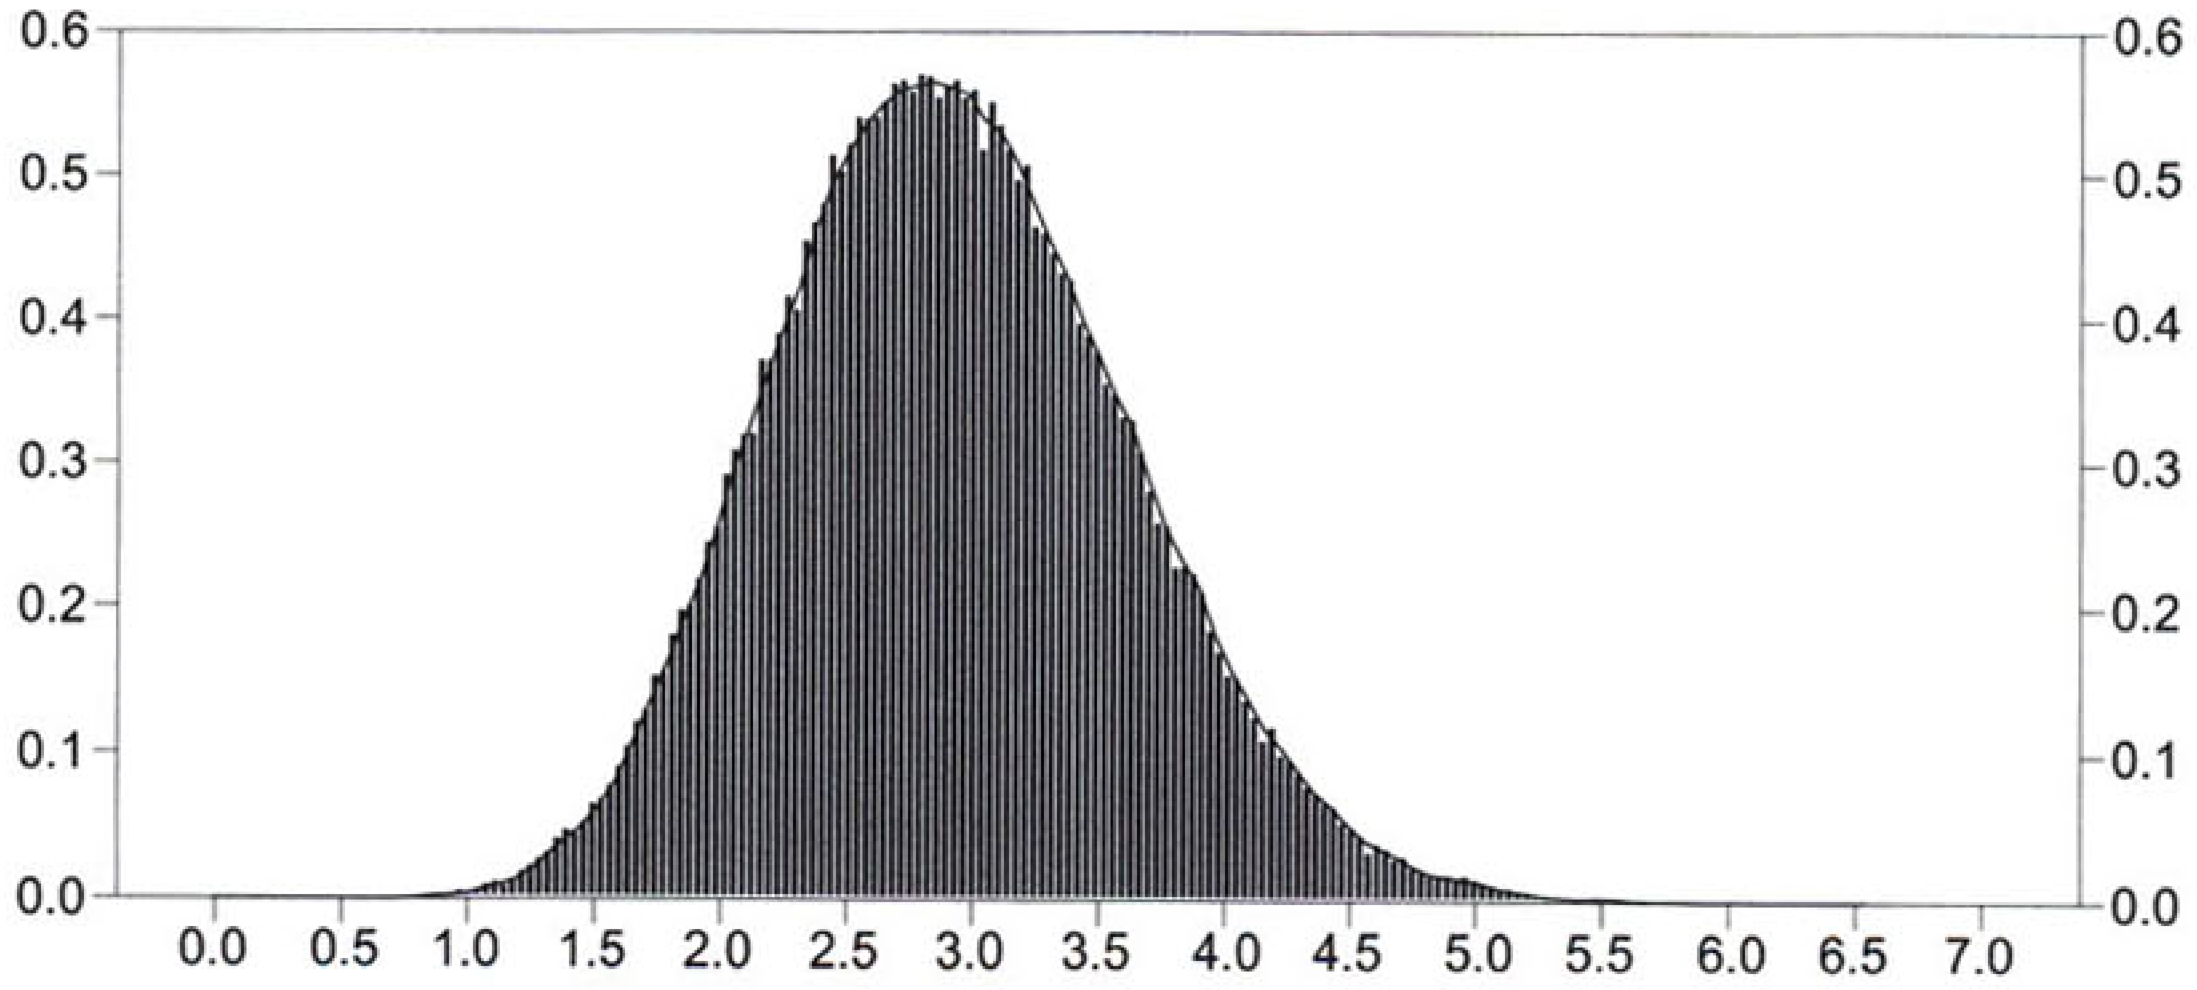

Note: Based on 100,000 random draws.

Exhibit B3

Mean Value of $OS^2 = (r'\Omega^{-1}r)$ for $r \sim t\left[\frac{(v-2)}{v}\Omega, v\right]$, $N = 9$, for Values of v in $3 \le v \le 100$

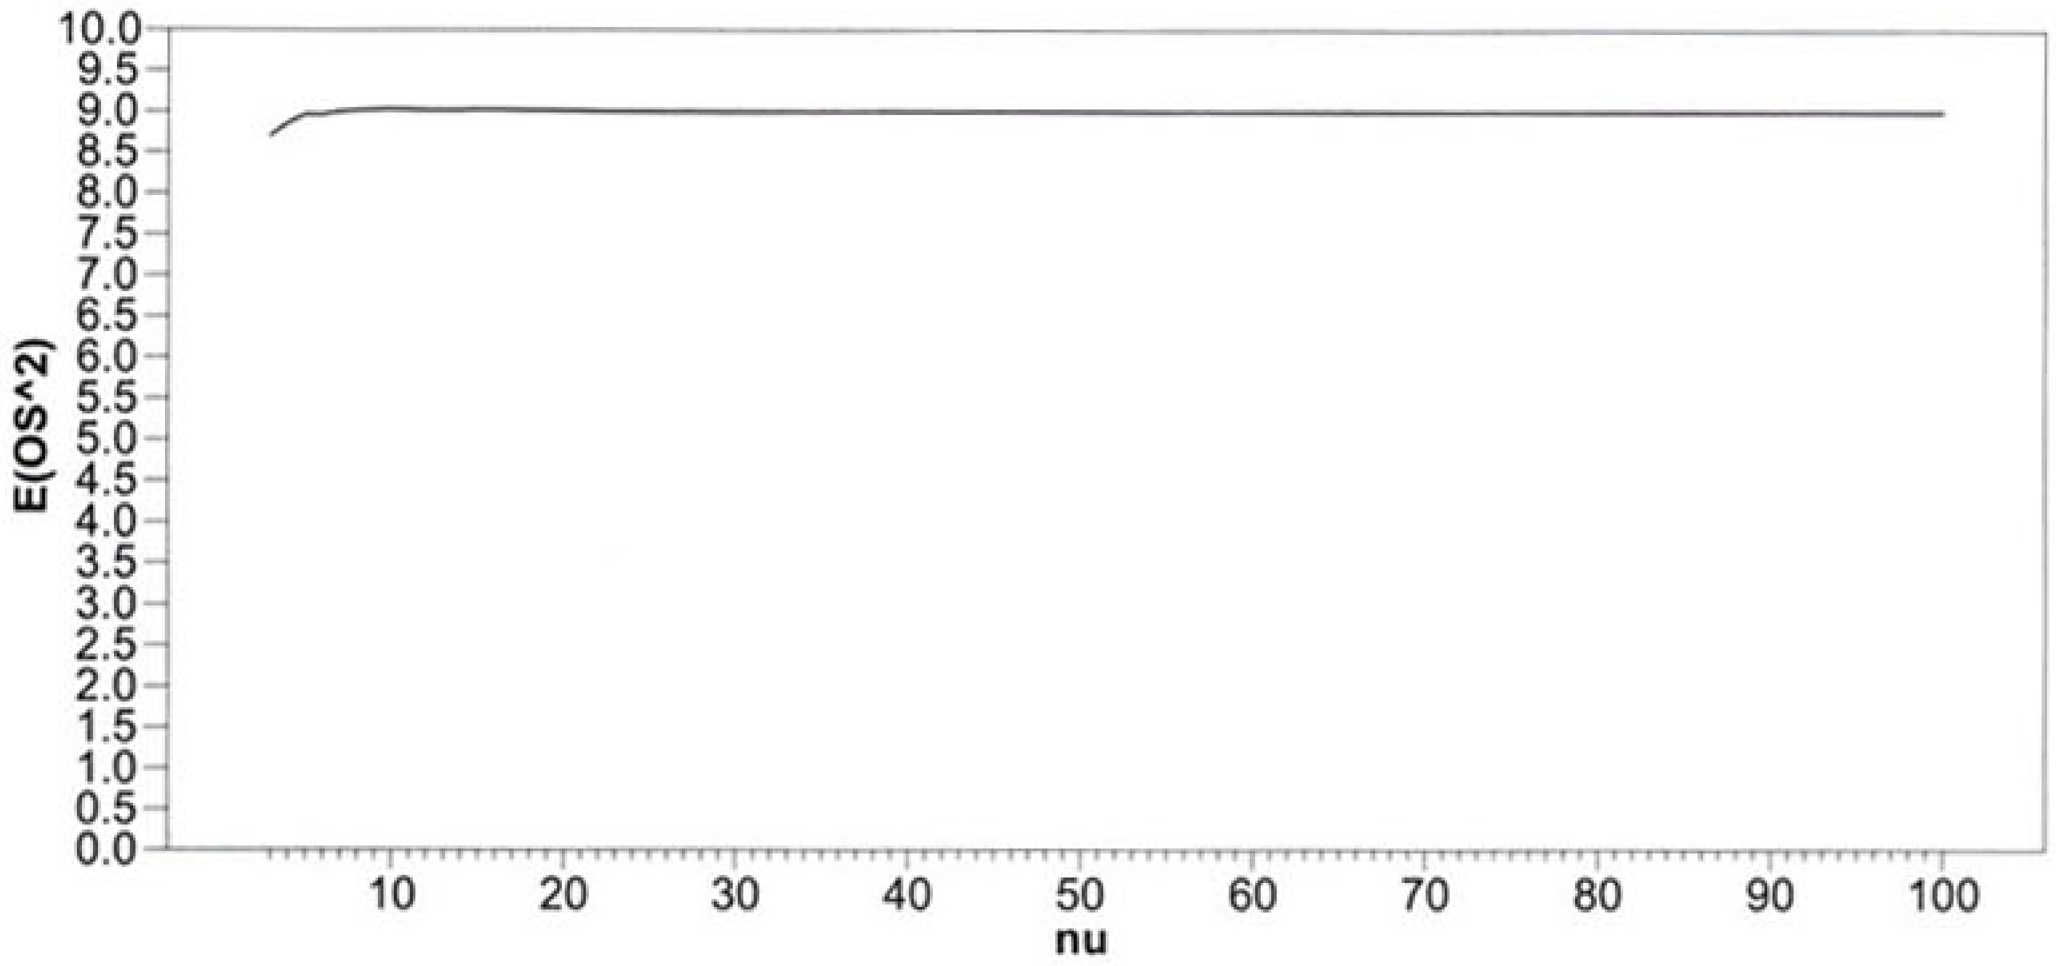

Note: Based on 100,000 random draws for each value of v.

Exhibit B4

Mean Value of $OS = \sqrt{(r'\Omega^{-1}r)}$, for $r \sim t\left[\frac{(v-2)}{v}\Omega, v\right]$, $N = 9$, for Values of v in $3 \le v \le 100$

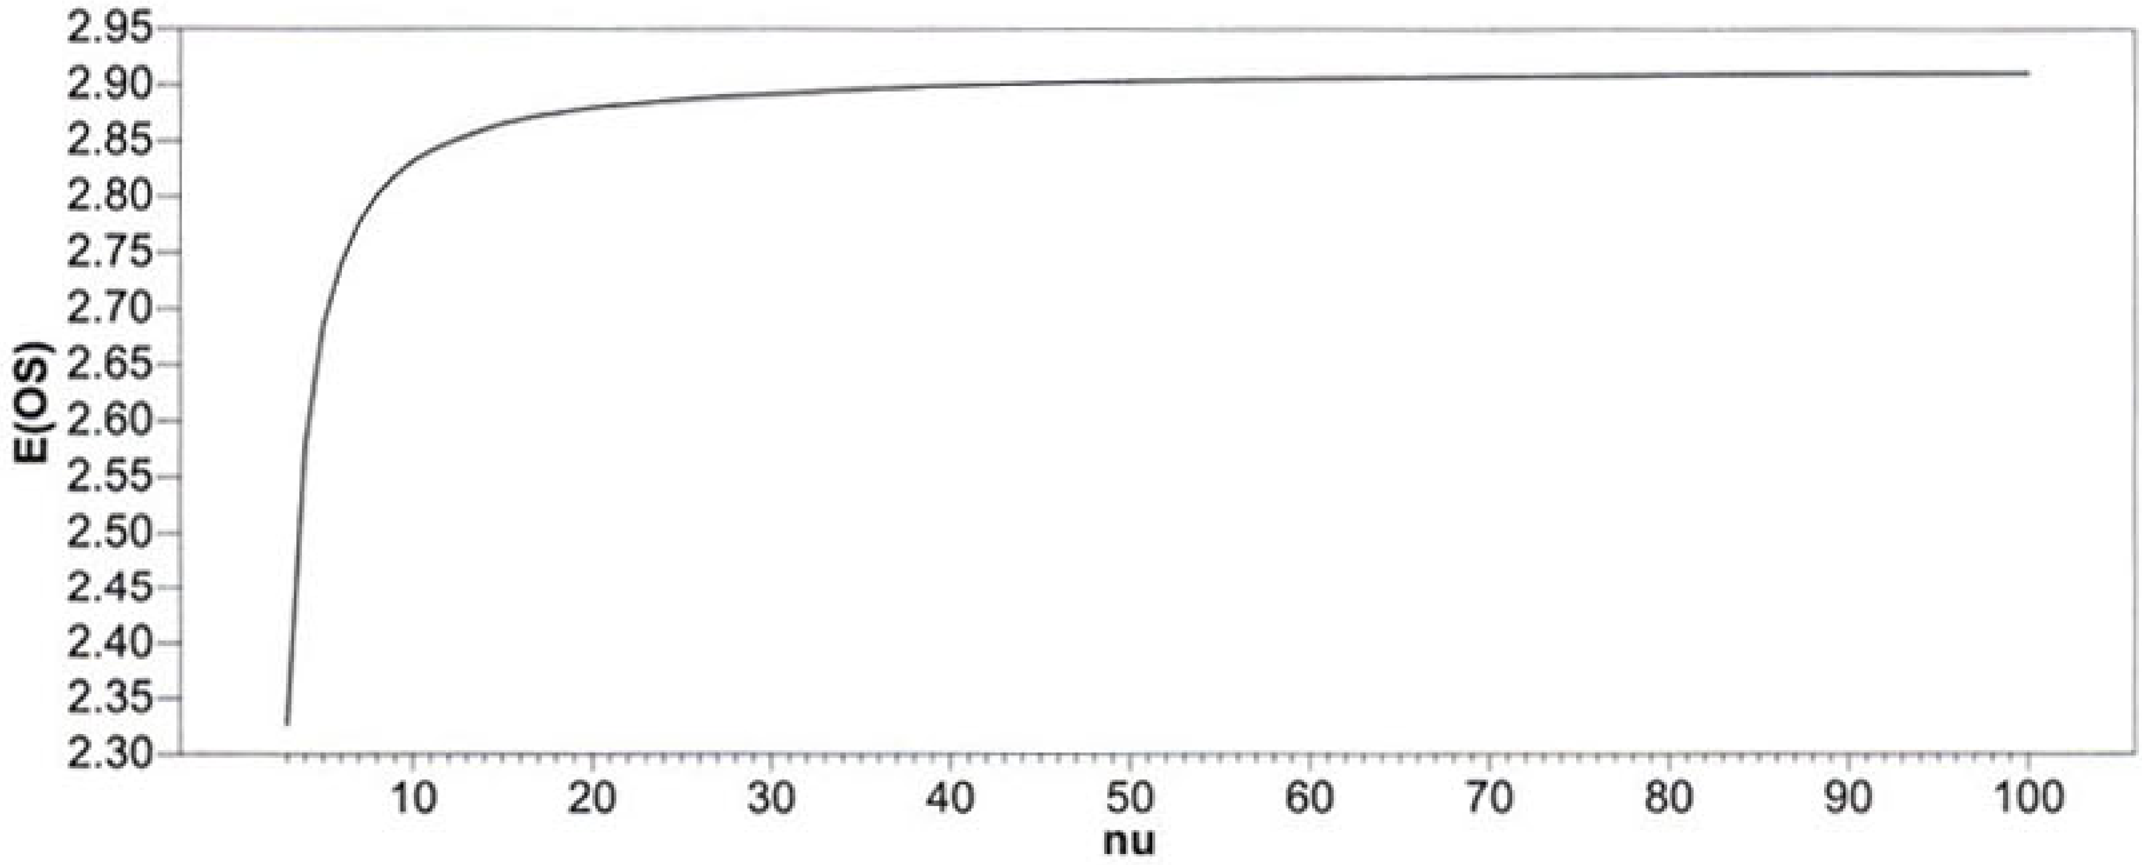

Note: Based on 100,000 random draws for each value of v.

Interestingly, it is quite symmetric, with only a very mild skew evident, and looks not unlike a normal distribution. This is, in fact, not surprising, because it can be shown (Fisher [1922]) that if a random variable X is distributed as χ^2 with N degrees of freedom (i.e., $X \sim \chi^2(N)$), then $\sqrt{2X}$ is approximately normally distributed with mean $\sqrt{2N-1}$ and unit variance (i.e., $\sqrt{2X} \sim Normal(\sqrt{2N-1}, 1)$). Given that for normally distributed returns $OS^2 \sim \chi^2(N)$, this implies that OS is approximately normally distributed with mean $\sqrt{(2N-1)/2}$ and variance 1/2. In particular, for $N = 9$, $E(OS) = \sqrt{(2N-1)/2} = 2.92$, which accords exactly with Exhibit B1, while the variance of 0.48 shown in Exhibit B1 is very close to the theoretical value of 0.5.

For $N > 6$, however, $\sqrt{(2N-1)/2}$ is less than 0.1 away from $\sqrt{N}$. Hence, for normally distributed returns and N assets, it seems reasonable to assume that the opportunity set of a portfolio is approximately normally distributed with mean $\sqrt{N}$ and variance 0.5, such that $OS \sim Normal(\sqrt{N}, 0.5)$.

EXHIBIT B5

Variance of $OS = \sqrt{(r'\Omega^{-1}r)}$ for $r \sim t\left[\frac{(\nu-2)}{\nu}\Omega, \nu\right]$, $N = 9$, for Values of ν in $3 \le \nu \le 100$

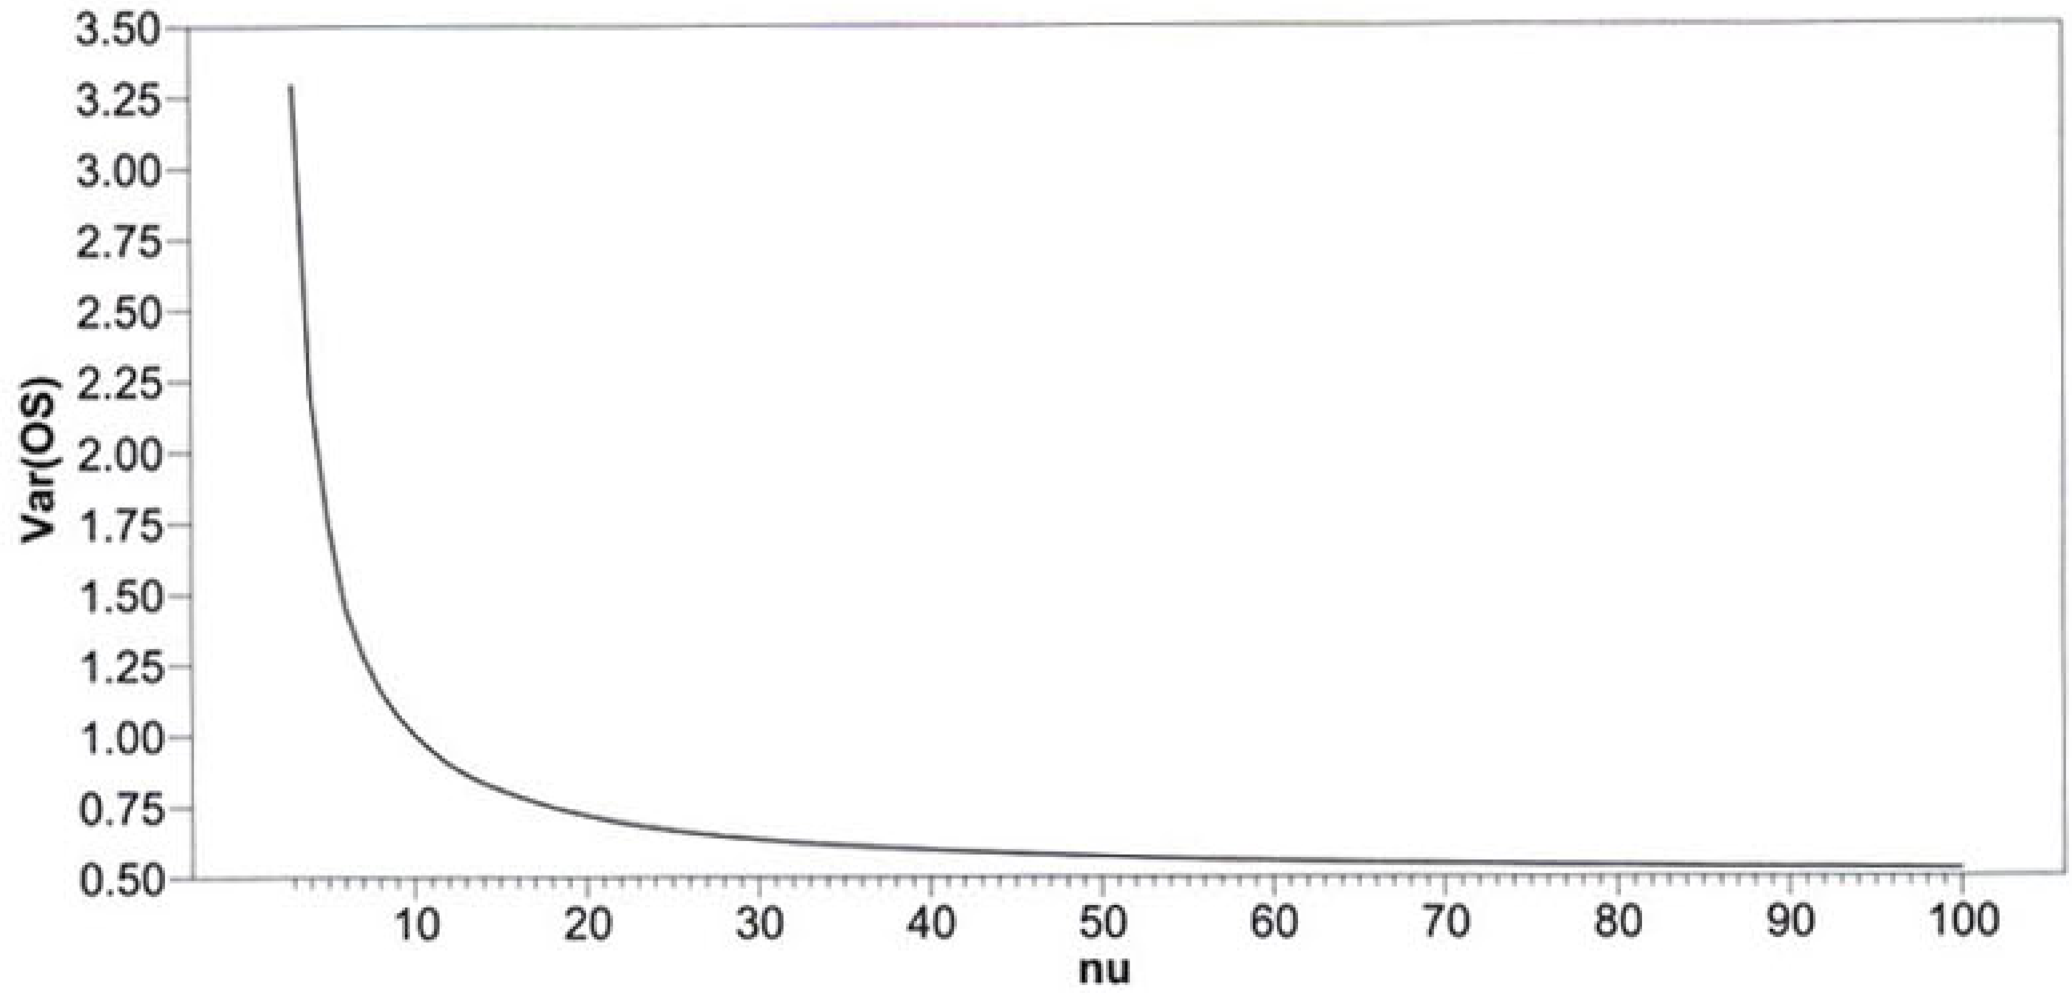

Note: Based on 100,000 random draws for each value of ν.

Assuming a Leptokurtic Distribution of Returns

The case of normally distributed returns is a limiting one, as asset returns are well known to be rather more leptokurtic (fat-tailed) than the normal distribution. While, as we have seen, clear analytic results can be derived under the assumption of normality, this is not so for the more general case of leptokurtic returns. Accordingly, we again investigated the distribution of OS using Monte Carlo methods, this time assuming a fat-tailed distribution of returns.

In particular, we constructed a set of Monte Carlo experiments in which the $N \times 1$ vector of returns was assumed to be drawn from a multivariate Student t distribution with parameters Σ and ν, $r \sim t(\Sigma, \nu)$. Σ is the so-called scale matrix, defined (for $\nu > 2$), such that $E(rr') = (\frac{\nu}{\nu-2})\,\Sigma = \Omega$ (i.e., $\Sigma = (\frac{\nu-2}{\nu})\Omega$). The parameter ν effectively determines the degree of leptokurtosis.[11] For low values of ν, the distribution will exhibit high excess kurtosis, while as $\nu \to \infty$, Student's t distribution becomes the normal distribution with zero excess kurtosis.[12]

We again assumed a nine-asset universe (i.e., $N = 9$) and estimated a 9×9 covariance matrix, Ω, using daily returns expressed in U.S. dollars for the Australian dollar, Canadian dollar, Euro, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, Swiss franc, and pound sterling over the period from January 1999 to March 2007.

We then drew r from a multivariate t distribution with covariance matrix Ω and $\nu = 3$ and calculated $OS = \sqrt{(r'\Omega^{-1}r)}$, repeating this for a total of 100,000 draws. The 100,000 computed values of OS were then taken as its empirical distribution for $\nu = 3$. We repeated this exercise for values of ν from 4 to 100.

Exhibit B3 shows the mean value of the squared opportunity set (i.e., $E(OS^2) = (r'\Omega^{-1}r)$) for the resulting empirical distributions for values of ν in the range $3 \le \nu \le 100$. This graph merely demonstrates by simulation methods, for $N = 9$, what we have shown analytically for all values of N, namely, that the result $E(OS^2) = N$ is distribution independent.[13]

Exhibit B4 shows the mean value of the opportunity set itself $\left(\text{i.e., } E(OS) = \sqrt{(r'\Omega^{-1}r)}\right)$ for the same set of empirical distributions. The first point to note about Exhibit B4 is that it illustrates, for the case of $N = 9$, the result that $E(OS) = E\sqrt{(r'\Omega^{-1}r)} < \sqrt{E(r'\Omega^{-1}r)} = N$, which follows from Jensen's inequality (i.e., that the mean value of a concave function of a random variable is less than the same concave function of the mean value of the same random variable). We showed above that, for normally distributed returns, $E(OS) = \sqrt{(2N-1)/2}$. For $N = 9$, this quantity is equal to 2.92. It is not surprising, therefore, that for high values of ν, as the t distribution approaches the normal distribution, $E(OS)$ approaches 2.92.

Exhibit B4 also shows that for distributions exhibiting greater kurtosis than the normal distribution, however, for lower values of ν, $E(OS)$ decreases in value. This again is unsurprising, because, as shown in Exhibit B5, the variance of OS rises as the distribution of returns becomes more leptokurtic and there is a higher probability of extreme events for any given covariance matrix of returns. Nevertheless, even for the most extreme fat-tailed distributions of returns (i.e., very low values of ν), $E(OS)$ falls at its lowest to a value of around 2.3, so that the approximation $E(OS) \approx \sqrt{N}$ (for $N = 9$) is still reasonable.

We showed in the text that the variance of OS for normally distributed returns is 0.5. This is illustrated in Exhibit B4, since for large values of ν the empirical variance of OS approaches this lower bound. For marked degrees of kurtosis (i.e., low values of ν), however, Exhibit B4 shows that for the case of $N = 9$, the variance of OS may be greater than 3.

ENDNOTES

[1]In fact, as discussed in Appendix B, for technical reasons (Jensen's inequality), the average value of OS will be close to, but slightly less than, $\sqrt{N}$.

[2]Suppose that S is the hindsight portfolio and U is the reference (market) portfolio. Let $\beta_s = \beta'_s = u'\Omega s / u'\Omega u$, then $q = s - \beta_s \cdot u$ will have $\theta_Q / \omega_Q = ROS$.

[3]We ran the analysis with and without the large returns. They are not materially different, although with exceptional large returns (550% in one month!), the largest value of $OS/\sqrt{N}$ is 1.69 not 1.29.

[4]If this analysis were done on a daily basis, the MOS would have been about 28 or 29 for the crash date.

[5]The Jarque-Bera test for normality yields p-values on the order of 10^{-12}.

[6]The risk forecast should have the same horizon as the returns. For example, weekly returns and an annual risk forecast will cause the measure of the opportunity set to be off by a factor of $\sqrt{52}$.

[7]The *trace* of a matrix is defined as the sum of its diagonal elements. This proof relies on two "tricks." The first is that a scalar (i.e., a simple number) can be viewed as a matrix with a single element, so that it is equal to its own trace. The second is that, for conformable vectors and matrix κ, μ, and Λ, $trace(\kappa' \Lambda \mu)$ is equal to $trace(\Lambda \mu \kappa')$. While this property of the *trace* operator is not immediately obvious, it is easily verified by means of simple examples.

[8]For example, in the second example, with $N = 13$ and the in-sample estimates of $\Omega^{-1}E(rr')$ and $\Omega^{-1}E(rr')$, we calculate a trace value of 13.043.

[9]The approximation is also due to Jensen's inequality, as discussed at length in Appendix B.

[10]With returns expressed in dollar terms, for example, including the dollar would introduce a singularity into the returns covariance matrix.

[11]For an $(N \times 1)$ zero-mean vector r with covariance matrix Ω, the density function of the multivariate $t(\Sigma, \mathrm{v})$ distribution is $\frac{\Gamma[(N+\nu)/2]}{\Gamma(\nu/2)\nu^{N/2}\pi^{N/2}|\Sigma|^{1/2}\left[1+(1/\nu)r'\Sigma^{-1}r\right]^{(\nu+N)/2}}$, where $\Gamma(\cdot)$ is the gamma function and $E(rr') = \left(\frac{\nu}{\nu-2}\right)\Sigma = \Omega$ for $\nu > 2$.

[12]The degree of excess kurtosis, defined for $\nu > 4$, is $6/(\nu - 4)$.

[13]The sharp-eyed reader will notice a very slight downturn in the values of $E(OS^2)$ for very low values of ν. In no case does $E(OS^2)$ fall below about 8.8, however, so the general result still seems valid.

REFERENCES

Clarke, R., H. De Silva, and S. Thorley. "Portfolio Constraints and the Fundamental Law of Active Management." *Financial Analysts Journal* (September/October 2002), pp. 48–66.

Fisher, R.A. "On the Interpretation of χ^2 from Contingency Tables and Calculation of P." *Journal of the Royal Statistical Society*, Series A, 85 (1922), pp. 87–94.

Grinold, R.C. "The Fundamental Law of Active Management." *Journal of Portfolio Management*, Vol. 15, No. 3 (Spring 1989), pp. 30–37.

——. "Attribution." *Journal of Portfolio Management*, Vol. 32, No. 2 (Winter 2006), pp. 9–24.

Grinold, R.C., and R.N. Kahn. *Active Portfolio Management*, 2nd ed. New York: McGraw-Hill, 2000.