Geometric Mean Maximization: *Expected, Observed, and Simulated Performance*

RAFAEL DE SANTIAGO AND JAVIER ESTRADA

RAFAEL DE SANTIAGO
is an associate professor of managerial decision sciences at IESE Business School in Barcelona, Spain.
rsantiago@iese.edu

JAVIER ESTRADA
is a professor of finance at IESE Business School in Barcelona, Spain.
jestrada@iese.edu

Academics and practitioners have developed a wide variety of approaches to optimize portfolios. In fact, although portfolio optimization traditionally referred to the maximization of the Sharpe ratio, nowadays the same expression can also be used to describe many alternative approaches, including optimization with higher moments, Bekaert et al. [1998]; full-scale optimization, Adler and Kritzman [2006]; and mean-semivariance optimization, Estrada [2008], to name but a few. These alternative approaches include geometric mean maximization, which is, together with the traditional criterion, the one we focus on in this article.

Sharpe ratio maximization implies selecting the portfolio with the highest risk-adjusted return, the latter defined as expected (excess) return per unit of volatility risk; geometric mean maximization, in turn, implies selecting the portfolio expected to grow at the fastest rate, therefore maximizing expected terminal wealth. Unfortunately, although both goals are desirable, selecting one implies, ex ante, giving up on the other. In other words, portfolios that aim to maximize risk-adjusted return are (typically very) different from those that aim to maximize the expected growth of the capital invested.

However, what is expected ex ante may be different from what actually happens ex post. In fact, when comparing the *observed* performance of portfolios that aim to maximize growth and those that aim to maximize risk-adjusted return, our results show that the former outperform the latter in terms of growth, and yet the former are *not* outperformed by the latter in terms of risk-adjusted return.

Furthermore, when comparing *simulated* performance, our results show that both criteria are likely to achieve their respective goals of maximizing growth or risk-adjusted return. Our results also show that despite its higher volatility, geometric mean maximization does not expose investors to substantially higher losses than does Sharpe ratio maximization. In fact, the former exposes investors to moderate losses not only at the end of, but also anywhere along, the holding period, and provides investors with far more upside potential than does Sharpe ratio maximization.

Our findings have important implications for portfolio managers. Ultimately, we find that portfolios resulting from geometric mean maximization, even (or perhaps particularly) when subject to diversification constraints, have very desirable characteristics. They are very likely to outperform in terms of growth, and provide substantial upside with rather limited downside, in both cases relative to portfolios resulting from Sharpe ratio maximization.

The criterion we focus on in this article has been variously referred to in the literature as the Kelly criterion, the growth optimal portfolio, the capital growth theory of investment, the geometric mean strategy, investment for the long run, or maximum expected log; here we will refer to it as *geometric mean maximization*, or GMM for short. And we will refer to the traditional criterion that aims to maximize risk-adjusted return as Sharpe ratio maximization, or SRM for short. Furthermore, the optimal portfolios that result from GMM and SRM are respectively referred to here as G (and G_C when constraints are added to GMM) and S.

The rest of the article is organized as follows. The second section briefly discusses the issue at stake. The third section discusses the implementation of the two optimization criteria evaluated in this article. The fourth section discusses the evidence on the expected, observed, and simulated performance of the portfolios generated by both optimization criteria. The final section provides an assessment. An appendix with Appendix A, Appendix B, and Appendix C, and methodology concludes the article.

THE ISSUE

The GMM criterion has a very long history—in fact, roughly as long as the history of the SRM criterion. The latter can be traced back to the seminal work of Markowitz [1952, 1959] and the complementary contributions of Treynor [1961], Sharpe [1964], Lintner [1965], and Mossin [1966]. The former, in turn, can be traced back to the seminal work of Kelly [1956] and Latane [1959]. Both Christensen [2005] and Poundstone [2005] provide thorough accounts of the origins and evolution of GMM and render an exhaustive literature review here unnecessary. Estrada [2010], from which the remainder of this section borrows heavily, also discusses the GMM criterion in detail.

Although Kelly [1956] focused on gambling and Latane [1959] on investing, both considered a set-up with many similarities; these include a gambler/investor making a large number of uncertain choices, a multiperiod framework, cumulative results, and the goal of maximizing expected terminal wealth (or, similarly, the expected growth of the capital invested). At the same time, the optimal strategies derived by both share many characteristics; these include that the allocations may be very aggressive, the capital invested may fluctuate widely over time, and betting/investing more (less) than suggested by the optimal strategy increases (decreases) risk and decreases expected terminal wealth.

Importantly, although SRM is a one-period framework, GMM is a *multiperiod* framework with *cumulative* results, which is consistent with the way most investors view and manage their portfolios. This distinction is critical because optimal decisions for a single period may be suboptimal in a multiperiod framework, and the relevant variable on which to focus when gains and losses are reinvested (the geometric mean) is different from the relevant variable when this is not the case (the arithmetic mean).[1]

Interestingly, although Latane [1959] proposed GMM as an alternative to Markowitz's framework, the latter has been one of the earliest and strongest supporters of this criterion. In fact, not only did he allocate the entire chapter VI of his pioneering book [Markowitz, 1959] to "Return in the Long Run," but he also added a "Note on Chapter VI" in a later edition. Markowitz [1976] reaffirmed his support for GMM.

Empirical research on the GMM criterion is rather scarce, and that is one of the voids this article aims to fill. Roll [1973] and Fama and MacBeth [1974] compare the G portfolio to the market portfolio and find that they are statistically indistinguishable.[2] Grauer [1981] finds that G portfolios are less diversified and have much higher expected return and volatility than S portfolios; Hunt [2005] finds similar results for the Australian market. Finally, using a sample of developed markets, emerging markets, and asset classes, Estrada [2010] confirms the relative (expected) characteristics of G and S portfolios already mentioned; he also finds that G portfolios are very likely to outperform S portfolios in terms of growth, and not likely to underperform in terms of risk-adjusted return, in both cases based on observed performance.

In short, then, this article aims to compare two portfolio optimization approaches, GMM and SRM; to assess the expected characteristics of the portfolios that stem from each criterion; and ultimately to evaluate the observed and simulated performance of those portfolios.

METHODOLOGY

Standard modern portfolio theory establishes that the expected return (μ_p) and variance (σ_p^2) of a portfolio are given by

$$\mu_p = \sum_{i=1}^{n} x_i \mu_i \tag{1}$$

$$\sigma_p^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij} \tag{2}$$

where x_i denotes the proportion of the portfolio invested in asset i; μ_i the expected return of asset i; σ_{ij} the covariance between assets i and j; and n the number of assets in the portfolio.

Maximizing risk-adjusted return when risk is measured with volatility amounts to maximizing a portfolio's Sharpe ratio (SR_p). This problem is formally given by

$$\text{Max}_{x_1, x_2, \ldots, x_n} \quad SR_p = \frac{\mu_p - R_f}{\sigma_p} = \frac{\sum_{i=1}^{n} x_i \mu_i - R_f}{\sqrt{\sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij}}} \tag{3}$$

$$\text{Subject to} \quad \sum_{i=1}^{n} x_i = 1 \quad \text{and} \quad x_i \geq 0 \text{ for all } i \tag{4}$$

where R_f denotes the risk-free rate and $x_i \geq 0$ the no short-selling constraint. This is the formal expression of the criterion referred to in this article as SRM; its resulting portfolio is referred to here as *S*. The solution of this problem is well known and available from a wide variety of optimization packages.

The maximization of a portfolio's geometric mean return can be implemented in more than one way. Ziemba [1972], Elton and Gruber [1974], Weide, Peterson, and Maier [1977], Bernstein and Wilkinson [1997], and Estrada [2010] all propose different algorithms to solve this problem. The method proposed here is easy to implement numerically and requires the same inputs as those needed for SRM. Following Estrada [2010], maximizing a portfolio's geometric mean return (GM_p) amounts to solving the problem formally given by

$$\begin{aligned} \text{Max}_{x_1, x_2, \ldots, x_n} & \; GM_p \\ & \approx \exp\left\{ \ln(1+\mu_p) - \frac{\sigma_p^2}{2(1+\mu_p)^2} \right\} - 1 \\ & \approx \exp\left\{ \ln\left(1+\sum_{i=1}^{n} x_i \mu_i\right) - \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} x_i x_j \sigma_{ij}}{2\left(1+\sum_{i=1}^{n} x_i \mu_i\right)^2} \right\} - 1 \end{aligned} \tag{5}$$

$$\text{Subject to} \quad \sum_{i=1}^{n} x_i = 1 \quad \text{and} \quad x_i \geq 0 \text{ for all } i \tag{6}$$

This is the formal expression of the criterion referred to in this article as GMM; its resulting portfolio is referred to here as *G* (or G_C if additional constraints are imposed; more on this later). Note that maximizing (5) is obviously the same as maximizing the expression inside the brackets. In fact, Markowitz [1959] suggests approximating the geometric mean of an asset precisely with the expression $\{\ln(1 + \mu) - \sigma^2/[2(1 + \mu)^2]\}$.

Finally, note that expression (5) highlights an important fact about the role that volatility plays in the GMM framework. In the SRM framework, volatility is undesirable because it is synonymous with risk; in the GMM framework, in turn, volatility is also undesirable, but for a different reason, namely, *because it lowers the geometric mean return*. In other words, in the GMM framework volatility is *not* ignored; it is detrimental because it lowers the rate of growth of the capital invested, thus ultimately lowering the expected terminal wealth.

EVIDENCE

We discuss in this section the main findings of our work. We focus first on comparing the expected characteristics of the *G* and *S* portfolios; then we assess the observed performance of these portfolios; and finally we evaluate their simulated behavior. Our sample consists of monthly returns for six assets classes, namely, U.S. stocks, EAFE stocks, emerging markets stocks, U.S. bonds, U.S. real estate, and gold. All returns are in dollars and account for capital gains/losses and dividends/coupons. The sample period varies by asset class but in all cases goes through December 2010. Exhibit A1 in the Appendix A describes the data in detail.

Expected Performance

Our first step consists of comparing the expected characteristics of the portfolios selected by GMM and SRM. In order to avoid drawing conclusions biased by particular temporal conditions, we optimize portfolios at three points in time: December 2000, December 2005, and December 2010. In all three cases, *S* portfolios follow from expressions (3)–(4) and *G* (and G_C) portfolios from expressions (5)–(6); also, in all cases, the inputs of the optimization problems (expected returns, variances, and covariances) are calculated on the basis of all the data available for each variable at the time of estimation. Exhibit 1 reports the relevant results.

EXHIBIT 1
Optimal Portfolios and Expected Performance

This exhibit shows optimal portfolios and some of their expected characteristics. The optimizations are performed at the end of December 2000, December 2005, and December 2010 based on all the data available at each point in time. S portfolios aim to maximize the Sharpe ratio and are obtained from expressions (3)–(4); G and G_C portfolios aim to maximize the geometric mean return and are obtained from expressions (5)–(6), with G_C constrained to have weights no larger than 47.5%. Panel A shows the weight of each asset in the optimal portfolios and panel B shows some of the portfolios' expected characteristics, including the number of assets in each (n), arithmetic (μ_p) and geometric (GM_p) mean return, volatility (σ_p), Sharpe ratio (SR_p), and the terminal value of $100 invested at GM_p after 10 (TV10), 20 (TV20), and 30 years (TV30). Mean returns, volatility, and Sharpe ratios in panel B are monthly magnitudes, unless indicated as annualized. The monthly risk-free rates used in the maximization of Sharpe ratios are 0.42% (December 2000), 0.36% (December 2005), and 0.28% (December 2010). The data is described in Exhibit A1 in the Appendix A.

	Dec/2000			Dec/2005			Dec/2010		
	S	G	G_C	S	G	G_C	S	G	G_C
Panel A: Weights (%)									
U.S. stocks	11.9	0.0	5.0	0.0	0.0	0.0	2.0	0.0	0.0
EAFE stocks	43.4	43.2	47.5	17.1	0.0	27.4	0.0	0.0	26.6
EM stocks	21.5	56.8	47.5	30.7	100.0	47.5	22.3	100.0	47.5
U.S. bonds	0.0	0.0	0.0	13.9	0.0	0.0	61.7	0.0	0.0
U.S. real estate	23.2	0.0	0.0	37.5	0.0	25.1	5.9	0.0	25.9
Gold	0.0	0.0	0.0	0.7	0.0	0.0	8.1	0.0	0.0
Panel B: Characteristics									
n	4	2	3	5	1	3	5	1	3
μ_p (%)	1.0	1.2	1.1	1.0	1.3	1.2	0.7	1.4	1.1
GM_p (%)	1.0	1.0	1.0	0.9	1.1	1.0	0.6	1.1	1.0
σ_p (%)	4.0	5.2	5.0	3.5	6.6	4.5	2.0	7.0	5.1
SR_p	0.157	0.142	0.146	0.181	0.149	0.175	0.195	0.154	0.166
Annualized GM_p (%)	12.2	13.0	12.9	11.8	14.4	13.3	8.0	14.2	12.6
Annualized σ_p (%)	13.7	18.2	17.2	12.2	23.0	15.7	6.9	24.1	17.7
TV10 ($)	315	340	338	306	385	350	216	377	328
TV20 ($)	994	1,154	1,142	935	1,485	1,224	468	1,425	1,073
TV30 ($)	3,136	3,921	3,858	2,860	5,724	4,284	1,014	5,378	3,516

We will focus for now on the G and S portfolios and come back to the G_C portfolios later. The characteristics of G and S at all three points in time are consistent with those previously reported in the literature. First, G portfolios are clearly less diversified than S portfolios; in fact, SRM never selects fewer than four assets, but GMM selects two assets in December 2000 and just one asset in December 2005 and December 2010. Second, the lower diversification of G portfolios makes them more volatile than S portfolios. Third, as expected by design, G portfolios are outperformed by S portfolios in terms of risk-adjusted return as measured by the Sharpe ratio. Fourth, and also as expected by design, G portfolios outperform S portfolios in terms of growth as measured by the geometric mean return. This in turn implies that G portfolios are expected to deliver a higher terminal capital, as the last three lines of the exhibit clearly show. In all cases, the differences in expected growth and terminal capital are substantial, particularly in the last two optimizations (December 2005 and December 2010).

Note that the (arithmetic and geometric) mean return, volatility, Sharpe ratio, and terminal capital reported are the expected characteristics of the G and S portfolios given the historical behavior of the assets they contain. But is this relative *expected* performance consistent with that actually *observed*? This is the issue we address in the next section.

Observed Performance

Exhibit 2 summarizes the results of the *observed* behavior of G and S; Exhibit B1 in the Appendix B

EXHIBIT 2
Observed Performance

This exhibit describes the observed performance of optimal portfolios defined as those that aim to maximize the Sharpe ratio (S) according to expressions (3)–(4) or the geometric mean return (G and G_C) according to expressions (5)–(6). The weights in G_C are constrained to be no larger than 47.5%. The figures summarize the performance of $100 invested in the optimal portfolios formed at the end of December 2000 (shown in Exhibit 1), passively held through the end of December 2010. The last column summarizes the performance of $100 passively invested in the world market (equity) portfolio over the same period. Performance measures include the arithmetic (μ_p) and geometric (GM_p) mean return, volatility (σ_p), semideviation with respect to 0 (Σ_p), beta with respect to the world market (β_p), lowest (Min) and highest (Max) return, Sharpe ratio (SR_p), and Sortino ratio ($N_p = \mu_p/\Sigma_p$), all expressed in monthly magnitudes, as well as the terminal value of the $100 investment (TV). The data is described in Exhibit A1 in the Appendix A.

	S	G	G_C	World
μ_p (%)	0.9	1.2	1.1	0.4
GM_p (%)	0.7	1.0	0.9	0.3
σ_p (%)	5.6	6.5	6.3	5.1
Σ_p (%)	4.0	4.4	4.3	3.7
β_p	1.1	1.2	1.2	1.0
Min (%)	−25.5	−25.6	−25.0	−19.8
Max (%)	17.5	15.9	15.4	11.9
SR_p	0.097	0.131	0.120	0.019
N_p	0.221	0.269	0.253	0.118
Annualized GM_p (%)	9.0	12.3	11.2	3.7
Annualized σ_p (%)	19.5	22.4	21.8	17.5
TV ($)	236	319	289	144

complements the analysis. (As before, we will come back to the G_C portfolio later.) The figures reported summarize the performance of a $100 investment in the G and S portfolios selected at the end of December 2000 (shown in Exhibit 1), passively held through the end of December 2010. The performance of a $100 passive investment in the world market (equity) portfolio over the same period is summarized in the last column simply for perspective.

Consistent with the expected characteristics discussed in the previous section, G is riskier than S regardless of whether risk is measured with the standard deviation (22.4% versus 19.5% in annualized terms), the semideviation, the beta, or the minimum monthly return. Also consistent with expectations, G outperforms S in terms of growth (12.3% versus 9.0% annualized geometric mean return) and terminal capital ($319 versus $236). The difference in annualized return, 330bps (basis points), is substantial and does not seem to come at a high price in terms of volatility. Furthermore, the terminal capital in G is 35% higher than that in S, and the $83 difference (=$319−$236) amounts to 83% of the initial investment ($100). Needless to say, these differences are substantial from an economic point of view.

Is it the case, then, that the higher compounding power of G relative to S is partially or fully offset by its higher volatility, thus producing a lower risk-adjusted return? No. As the exhibit shows, the Sharpe ratio of the G portfolio (0.131) is actually higher than that of the S portfolio (0.097). Although the difference is not statistically significant, at the very least these figures show that G is not outperformed by S in terms of risk-adjusted return.[3] These results also hold if risk is measured with the semideviation and risk-adjusted return with the Sortino ratio.

To summarize, the observed behavior of the G and S portfolios over the January 2001-December 2010 period is partly as expected and partly somewhat unexpected. As expected, G is more volatile, grows more rapidly, and leads to a higher terminal capital than S. But, perhaps surprisingly, G is not outperformed by S in terms of risk-adjusted return.

Observed Performance—Constrained GMM

The analysis in the previous section could be criticized on at least two grounds. First, it could be argued that even taking the results at face value, most investors would be reluctant to hold portfolios as concentrated as those selected by the GMM criterion. And second, it could be argued that the results reported are relevant only in the two paths examined (one for each criterion). We consider the first issue in this section and the second issue in the next one.

As shown in Exhibit 1, the GMM criterion selects two assets in December 2000 and only one in December 2005 and December 2010. Needless to say, such a low degree of diversification would be difficult to digest for most investors, even accounting for the fact that the assets considered are diversified *within* each asset class. And yet, this concentration should not lead to the rejection of GMM; rather, it should lead to the specification of the necessary diversification constraints.

To that purpose, we re-optimize portfolios on the same three dates as before, but this time constraining GMM to invest no more than 47.5% of the portfolio in any given asset. With this constraint we guarantee, first, that the resulting portfolio will have at least three assets; and second, that none of these assets will have a weight lower than a meaningful 5%. The expected and observed behavior of the resulting constrained G portfolios is reported in the columns labeled G_C in Exhibits 1 and 2; Exhibit B2 in the Appendix B complements the analysis.

Exhibit 1 shows that G_C portfolios have three assets at all three points in time. It also shows that, even with these constraints, the portfolios selected by GMM are still expected to outperform those selected by SRM in terms of growth and terminal capital (and to be outperformed in terms of risk-adjusted return).

Exhibit 2, which summarizes observed performance over the January 2001-December 2010 period, shows that G_C outperforms S in terms of growth (11.2% versus 9.0% annualized geometric mean return) and terminal capital ($289 versus $236). In fact, the difference in annualized return (220bps) remains substantial and comes at a low price in terms of volatility (21.8% versus 19.5% in annualized terms). Furthermore, the terminal capital in G_C is 22% higher than that in S, and the $53 difference (=$289−$236) amounts to 53% of the initial investment ($100). In short, the differences in growth and terminal capital remain substantial even after adding diversification constraints to GMM.

Importantly, the Sharpe ratio of the G_C portfolio (0.120) is higher than that of the S portfolio (0.097), though not significantly so from a statistical point of view. In other words, although G_C outperforms S in terms of growth and terminal capital, it is not outperformed by S in terms of risk-adjusted return. These results also hold if risk is measured with the semideviation and risk-adjusted return with the Sortino ratio.

Finally, note that the observed performance of the S, G, and G_C portfolios discussed is that of a buy-and-hold strategy; that is, $100 is invested in each portfolio at the end of December 2000 and passively held through the end of December 2010. Exhibit A2 in the Appendix A considers an alternative scenario in which portfolios are rebalanced halfway through the 10-year observation period. More precisely, $100 is invested in the optimal S, G, and G_C at the end of December 2000 (shown in Exhibit 1) and passively held through the end of December 2005; the capital accumulated in these portfolios is then reallocated to the optimal S, G, and G_C estimated at that point in time (also shown in Exhibit 1); and these portfolios are passively held through the end of December 2010. As the exhibit shows, this rebalancing halfway into the observation period does not substantially affect any of the results discussed.

Simulated Performance

The evidence on the observed performance of S, G, and G_C portfolios suggests that GMM should at the very least be considered a serious alternative to SRM. The fact that G and G_C outperform S in terms of growth and terminal capital, but are not outperformed by S in terms of risk-adjusted return, underscores the plausibility of GMM. However, although this evidence is based on observed performance, it is also based on the "one sample of history" (as Paul Samuelson would say) that we have actually observed. For this reason, we explore in this section the behavior of S, G, and G_C portfolios in thousands of other scenarios that could have happened.

The methodology behind our simulations, very briefly, is as follows. (Technical details are discussed in Appendix C.) First, we estimate the mean returns, volatilities, and correlations of the six assets classes in our sample with all the information available at the end of December 2010. Then we use that information to determine the S, G, and G_C portfolios at that point in time. These two steps result in the three optimal portfolios shown in the last three columns of Exhibit 1. We then simulate 10,000 paths for each S, G, and G_C portfolio over the 10-year (120-month) period between January 2011 and December 2020, thus running 10,000 horse races. Finally, we calculate several performance measures to summarize the results of these 10,000 horse races. Our main results are shown in Exhibits 3 and 4.

Panel A of Exhibit 3 summarizes the average simulated performance of the S, G, and G_C portfolios across the 10,000 scenarios considered for the January 2011-December 2020 period. To illustrate, for each of the 10,000 paths for S over the 120-month simulation period, we calculate its geometric mean monthly return; the average of those 10,000 figures is 0.7%, and the respective figures for G and G_C are 1.1% and 1.0%. The interpretation of the rest of the figures in this panel is similar.

As panel A shows, then, the differences in growth and terminal capital in the S, G, and G_C portfolios are remarkable. On average, G and G_C portfolios respectively outperform S portfolios by 540bps and 360bps a year, as indicated by their annualized geometric mean returns of 14.2%, 12.4%, and 8.8%. These differences imply, again on average, a terminal capital in G portfolios ($501) over twice as high as that in S portfolios ($250), and 58% higher in G_C portfolios ($395) than in S portfolios.

Panel A also shows that G and G_C portfolios are on average more volatile than S portfolios, as indicated by their respective annualized volatilities of 25.2%, 20.8%, and 10.0%. This higher volatility imposes a heavy drag on risk-adjusted return, leading G and G_C portfolios to underperform S portfolios as indicated by their respective Sharpe ratios of 0.186, 0.190, and 0.254. These results also hold if risk is measured with the semideviation and risk-adjusted return with the Sortino ratio.

Panel B shows the proportion of the 10,000 horse races in which G and G_C portfolios beat S portfolios in terms of growth (hence terminal capital) and risk-adjusted return, the latter measured both with the Sharpe ratio and the Sortino ratio. As the panel shows, G (G_C) portfolios produce higher growth than S portfolios 82.9% (81.1%) of the time. Conversely, G (G_C) portfolios produce higher Sharpe ratios than S portfolios only 15.8% (16.6%) of the time, and higher Sortino ratios 16.8% (17.3%) of the time. In short, across the 10,000 simulated scenarios, both GMM and SRM seem to achieve their respective goals most of the time.

Panel C focuses on the capital accumulated *at the end* of the 10-year simulation period. As already mentioned, the average terminal capital in the S, G, and G_C portfolios across the 10,000 scenarios is $250, $501, and $395, thus implying a substantial edge for GMM. The spread between the worst scenario and the best scenario is, as expected, larger for G ($19 and $9,306) and G_C

EXHIBIT 3
Simulated Performance

This exhibit shows results from our 10,000 simulations each over a 120-month period. S portfolios aim to maximize the Sharpe ratio and are obtained from expressions (3)–(4); G and G_C portfolios aim to maximize the geometric mean return and are obtained from expressions (5)–(6), with G_C constrained to have weights no larger than 47.5%. Panel A shows averages across 10,000 paths for the arithmetic (μ_p) and geometric (GM_p) mean return, volatility (σ_p), semideviation with respect to 0 (Σ_p), lowest (Min) and highest (Max) return, Sharpe ratio (SR_p), and Sortino ratio ($N_p = \mu_p/\Sigma_p$), all expressed in monthly magnitudes, as well as for the terminal value of the $100 investment (TV). Panel B shows the percentage of the 10,000 paths in which G and G_C beat S in the dimensions indicated. Panel C summarizes information about TVs, including the average (Avg), lowest (Min), and highest (Max) values across the 10,000 paths, as well as the average value in the quartile and decile with lowest terminal capital (Q1 and D1) and the quartile and decile with highest terminal capital (Q10 and D10).

	S	G	G_C
Panel A			
μ_p (%)	0.7	1.4	1.1
GM_p (%)	0.7	1.1	1.0
σ_p (%)	2.9	7.3	6.0
Σ_p (%)	1.7	4.4	3.6
Min (%)	−7.2	−17.5	−14.5
Max (%)	8.5	20.2	16.7
SR_p	0.254	0.186	0.190
N_p	0.457	0.322	0.329
Annualized GM_p (%)	8.8	14.2	12.4
Annualized σ_p (%)	10.0	25.2	20.8
TV ($)	250	501	395
Panel B (%)			
GM_p		82.9	81.1
SR_p		15.8	16.6
N_p		16.8	17.3
Panel C ($)			
Avg	250	501	395
Min	81	19	26
Max	2,255	9,306	5,034
Avg Q1	152	141	142
Avg Q4	399	1,099	791
Avg D1	133	94	102
Avg D10	510	1,569	1,081

(\$26 and \$5,034) portfolios than for S portfolios (\$81 and \$2,255).

However, focusing on only two scenarios (the best and the worst) out of 10,000 may be misleading. For this reason, we take the terminal capital in S portfolios for the 10,000 scenarios, rank them from the lowest to the highest, and calculate the average terminal capital for the top and bottom quartiles (and deciles); we then do the same for the G and G_C portfolios.

Interestingly, as panel C shows, the average terminal capital in the worst quartile is not much lower for G (\$141) and G_C (\$142) portfolios than for S portfolios (\$152). At the same time, the average terminal capital in the best quartile is *much* higher for G (\$1,099) and G_C (\$791) portfolios than for S portfolios (\$399). These results suggest the existence of an important asymmetry in upside and downside potential when investing in G and G_C portfolios as opposed to in S portfolios. Put differently, although in the "bad" scenarios an investor would be expected to fare somewhat worse by investing in G and G_C than in S, in the "good" scenarios the investor would be expected to fare *much* better.

Importantly, investors are typically concerned about the probability and magnitude of potential losses. For this reason, we explore the proportion of the 10,000 scenarios in which S, G, and G_C portfolios are under \$100 (the initial capital) *at the end* of the 10-year simulation period. Panel A of Exhibit 4 shows the proportion of paths that end with different levels of losses for all S, G, and G_C portfolios.

At all the levels of loss considered, G and G_C portfolios end up with a higher proportion of paths under \$100 than S portfolios. However, the proportion of paths in which G and G_C portfolios end with losses is very low. In only 4.0% (2.9%) of the scenarios considered, G (G_C) portfolios end with losses higher than 10%; and in only 2.8% (1.9%) of the scenarios considered, G (G_C) portfolios end with losses higher than 20%. In short, then, G and G_C portfolios are more likely than S portfolios to end a 10-year holding period with losses, but the probability of this happening is very low.

That being said, not all investors focus only on what happens *at the end* of any given holding period. As argued by Kritzman and Rich [2002], many investors do (or should) care about what happens *throughout* the holding period. In other words, it is important to assess the likelihood and magnitude of losses not just at the end of, but also *anywhere along* any given holding period.

Exhibit 4

Simulated Performance—Downside Potential

This exhibit shows results from our 10,000 simulations, each over a 120-month period, focusing on losses. S portfolios aim to maximize the Sharpe ratio and are obtained from expressions (3)–(4); G and G_C portfolios aim to maximize the geometric mean return and are obtained from expressions (5)–(6), with G_C constrained to have weights no larger than 47.5%. Panel A focuses on losses *at the end* of each path (10,000 months) and shows the percentage of the 10,000 paths that accumulate different levels of losses. Panel B focuses on losses *anywhere along* each path (1.2 million months) and shows the percentage of the 10,000 paths that accumulate different levels of losses.

	S	G	G_c
Panel A (%)			
Loss > 0%	0.2	5.2	4.2
Loss > 10%	0.1	4.0	2.9
Loss > 20%	0.0	2.8	1.9
Loss > 30%	0.0	1.9	1.1
Panel B (%)			
Loss > 0%	5.8	15.4	13.7
Loss > 10%	0.5	9.9	7.5
Loss > 20%	0.0	5.9	3.8
Loss > 30%	0.0	3.3	1.7

Note that for each criterion (S, G, and G_C) we simulate 10,000 paths of 120 months each, which amounts to a total of 1.2 million simulated months per criterion. Panel B of Exhibit 4 shows the proportion of these 1.2 million months in which S, G, and G_C portfolios are under \$100. As in panel A, it remains the case that at all the levels of losses considered, G and G_C portfolios spend more months under \$100 than S portfolios. But, also as before, the proportion of months with losses is rather low in all cases. Note that G (G_C) portfolios accumulate losses higher than 10% less than 10% (8%) of the time, and losses higher than 20% less than 6% (4%) of the time. In other words, even when considering not just what happens at the end of, but anywhere along the 10,000 simulated paths, it is still the case that G and G_C portfolios do not expose investors to much higher losses than do S portfolios.

To summarize, panel C of Exhibit 3 shows that GMM exposes investors to much higher upside potential than does SRM. The same panel and Exhibit 4, in turn, show that despite its high volatility, GMM does not expose investors to a considerable downside potential. These results combined suggest that GMM provides both a substantial upside and a rather limited downside, which should make it an attractive criterion for investors and portfolio managers.

AN ASSESSMENT

Portfolio optimization has become a crowded field, with many competing approaches in which Sharpe ratio maximization (SRM) remains the standard criterion. The results we discuss in this article, based on expected, observed, and simulated performance, suggest that geometric mean maximization (GMM) is a plausible criterion that should be seriously considered by both academics and practitioners.

There is no denying that GMM typically selects portfolios (G) that are much less diversified and much more volatile than those selected by SRM (S). And yet that shortcoming may easily be overcome by imposing the necessary diversification constraints. Our results show that diversification-constrained GMM selects portfolios (G_C) that retain most of the desirable characteristics of the portfolios selected by unconstrained GMM.

Our results also show that over the January 2001–December 2010 period, both G and G_C portfolios outperformed S portfolios in terms of growth, as measured by the geometric mean return and terminal capital, and yet did not underperform in terms of risk-adjusted return, as measured by the Sharpe and Sortino ratios. In fact, the observed (*out-of-sample*) annualized return differential with respect to S was a remarkable 330bps in the case of G, and a substantial 220bps in the case of G_C.

Our simulations further strengthened the appeal of GMM. In the 10,000 paths we simulated for each criterion over a 10-year holding period, G (G_C) portfolios outperformed S portfolios by 540bps (360bps) a year, thus producing much higher levels of terminal capital. In fact, top-quartile terminal capital in G (G_C) portfolios was more than 2.7 times (almost 2 times) higher than that in S portfolios. These differences would obviously be even larger in holding periods longer than 10 years.

Interestingly, the much higher upside potential of G and G_C portfolios was not offset by much higher downside potential. Our simulations show that bottom-quartile terminal capital in G and G_C portfolios was roughly just 7% lower than that in S portfolios. Furthermore, although G and G_C portfolios were more likely to be underwater than S portfolios, both during and at the end of the holding period, the probability of being underwater was rather low. In our simulations, G (G_C) portfolios accumulated losses higher than 10% less than 10% (8%) of the time, and losses higher than 20% less than 6% (4%) of the time, in both cases considering performance not just at the end of, but anywhere along the holding period. In other words, G and G_C portfolios are not likely to expose investors to much higher losses than S portfolios.

What kind of investors would benefit the most from GMM? Estrada [2010] argues that GMM is more attractive 1) the lower the degree of risk aversion; 2) the longer the holding period; and 3) the more certain the holding period. Obviously, the less risk averse an investor, the better he can tolerate the high volatility of the portfolios selected by this criterion. And naturally, the longer the holding period, the more time GMM has to deliver its higher expected growth; in the short term, anything can happen, and luck may play an important role (whose impact would be expected to decrease as the holding period increases).

As for the certainty of the holding period, if an investor's portfolio is not substantial and is likely to be used to take care of unforeseen contingencies, then the likelihood of having to liquidate it earlier than expected may be high. In these circumstances, an investor may intend to take the long view but may be forced to exit the strategy before it has time to deliver its expected higher growth. Similarly, a portfolio manager may want to take the long view, but the investors in his fund may be intolerant to suffering short-term losses and likely to exit the fund when these materialize. In short, the higher the probability to *remain* invested for the *long* term, the more attractive GMM becomes.

Long-term investors, portfolio managers whose funds attract long-term investors, and hedge funds (which typically impose lock-up periods that force investors to take the long view) may benefit the most from GMM. Relative to the widely accepted SRM criterion, then, GMM provides a much higher upside potential with a rather limited downside potential, and that should make it a plausible choice for investors and portfolio managers.

Appendix A

Exhibit A1
Data and Summary Statistics

This exhibit shows, for the series of monthly returns, the arithmetic (μ) and geometric (GM) mean return, standard deviation (σ), beta with respect to the world market (β), index of standardized skewness (SSkw), and index of standardized kurtosis (SKrt) for the six asset classes in the sample and for the world market, all of them calculated between the beginning (Start) and the end (December 2010) of each asset's sample period. The returns of U.S. stocks are summarized by the S&P total return index (from Global Financial Data). The returns of EAFE (Europe, Australasia, and the Far East) stocks and EM (Emerging Markets) stocks are summarized by MSCI total return indices. The returns of U.S. bonds are summarized by the 10-year government bond total return index (from Global Financial Data), and those of U.S. real estate by the FTSE NAREIT (All REITs) total return index. The return of gold is based on its New York price ($/ounce). The world market is summarized by the MSCI All Country World index. All returns are in dollars and account for capital gains/losses and dividends/coupons.

Asset Class	μ(%)	GM(%)	σ(%)	β	SSkw	SKrt	Start
U.S. stocks	0.9	0.8	5.2	0.84	5.0	73.3	Jan/1900
EAFE stocks	0.9	0.8	5.0	1.07	–2.8	5.1	Jan/1970
EM stocks	1.4	1.1	7.0	1.16	–4.7	6.0	Jan/1988
U.S. bonds	0.4	0.4	1.7	–0.03	12.6	49.2	Jan/1900
U.S. real estate	0.9	0.8	5.2	0.60	–3.7	34.1	Jan/1972
Gold	0.5	0.4	4.6	0.05	16.4	53.4	Jan/1940
World (Stocks)	0.7	0.6	4.5	1.00	–4.5	5.5	Jan/1988

Exhibit A2
Observed Performance—With Rebalancing

This exhibit describes the observed performance of optimal portfolios defined as those that aim to maximize the Sharpe ratio (S) according to expressions (3)–(4) or mean compound return (G and G_C) according to expressions (5)–(6). The weights in G_C are constrained to be no larger than 47.5%. The figures summarize the performance of $100 invested in the optimal portfolios formed at the end of December 2000 (shown in Exhibit 1); passively held through the end of December 2005; rebalanced to the optimal portfolios formed at the end of December 2005 (shown in Exhibit 1); and passively held through the end of December 2010. The last column summarizes the performance of $100 passively invested in the world market (equity) portfolio. Performance measures include the arithmetic (μ_p) and geometric (GM_p) mean return, volatility (σ_p), semideviation with respect to 0 (Σ_p), beta with respect to the world market (β_p), lowest (Min) and highest (Max) return, Sharpe ratio (SR_p), and Sortino ratio ($N_p = \mu_p/\Sigma_p$), all expressed in monthly magnitudes, as well as the terminal value of the $100 investment (TV). The data is described in Exhibit A1 in the Appendix A.

	S	G	G_C	World
μ_p (%)	0.9	1.3	1.0	0.4
GM_p (%)	0.7	1.1	0.8	0.3
σ_p (%)	5.0	6.8	6.2	5.1
Σ_p (%)	3.6	4.6	4.3	3.7
β_p	0.9	1.2	1.2	1.0
Min (%)	–23.3	–27.4	–26.3	–19.8
Max (%)	13.7	17.1	17.6	11.9
SR_p	0.106	0.142	0.110	0.019
N_p	0.244	0.284	0.236	0.118
Annualized GM_p (%)	9.3	13.6	10.3	3.7
Annualized σ_p (%)	17.5	23.5	21.5	17.5
TV ($)	243	357	268	144

APPENDIX B

EXHIBIT B1
Observed Performance

This exhibit shows the performance of $100 invested at the end of December 2000, passively held through the end of December 2010, in two optimal portfolios, one selected by SRM (S) and the other selected by GMM (G). It also shows the performance of $100 passively invested in the world market portfolio (W). Related performance figures are shown in Exhibit 2.

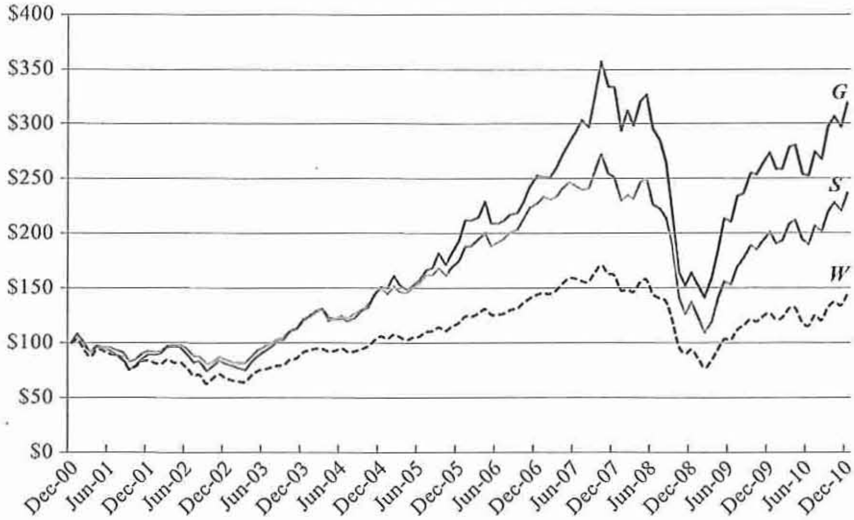

EXHIBIT B2
Observed Performance—Constrained GMM

This exhibit shows the performance of $100 invested at the end of December 2000, passively held through the end of December 2010, in three optimal portfolios, one selected by SRM (S), one selected by GMM (G), and one selected by GMM constrained to have weights no larger than 47.5% (G_C). It also shows the performance of $100 passively invested in the world market portfolio (W). Related performance figures are shown in Exhibit 2.

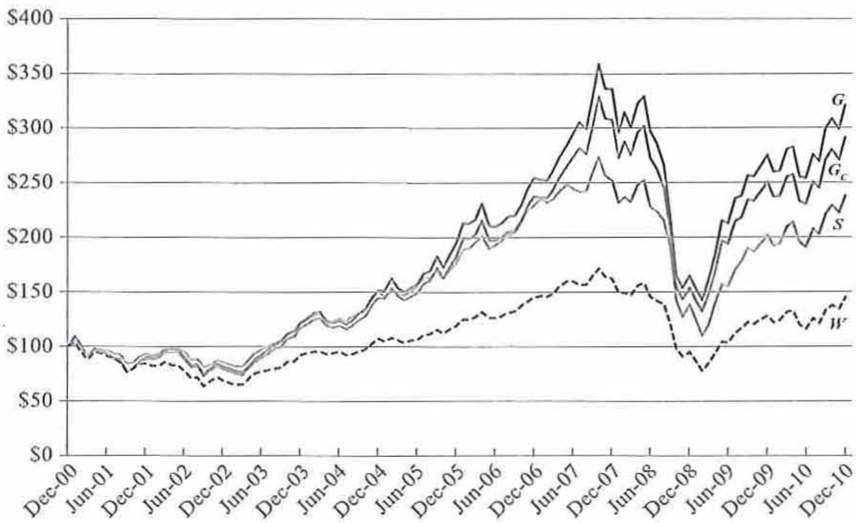

APPENDIX C

SIMULATIONS—METHODOLOGY

We describe in this section in more detail the methodology behind our simulations, very briefly discussed before in the "Simulated Performance" section. These simulations generate 10,000 (out-of-sample) scenarios for each S, G, and S portfolio; Exhibits 3 and 4 in the text report several performance measures related to these 10,000 horse races.

On January 1, 2011, an initial capital of $100 is allocated to each S, G, and G_c portfolio, and the evolution of each $100 is simulated over the subsequent 120 months. Note that simulated data correspond to the "future" in the sense of corresponding to a period after which our sample ends. Our simulations use all the information available as of the end of December 2010, and generate potential paths of asset returns for the following 10 years (120 months).

Importantly, the evolution of each portfolio cannot be simulated independently, because the asset classes we focus on are correlated.[4] For this reason, we need to simulate separately the evolution of each asset class, including their correlations, and then aggregate the results according to the composition of each optimal portfolio (shown in the last three columns of panel A in Exhibit 1).[5]

We thus start by estimating mean returns, volatilities, and covariances for our six asset classes following the RiskMetrics Exponentially Weighted Moving Average (EWMA) forecasting approach.[6] We do so with some minor modifications, such as not assuming zero mean returns; instead, we use the actual means estimated from the full sample available for each asset class. Then, for any two assets i and j, we estimate volatilities (σ_i) and covariances (σ_{ij}) with the expressions

$$\sigma_i = \left\{ (1-\lambda) \cdot \sum_{t=1}^{120} \lambda^{t-1} \cdot (r_{it} - m_i)^2 \right\}^{1/2}$$

$$\sigma_{ij} = (1-\lambda) \cdot \sum_{t=1}^{120} \lambda^{t-1} \cdot (r_{it} - m_i)(r_{jt} - m_j)$$

where r_{it} denotes the return of asset i in month t, m_i denotes the mean return of asset i, and λ is a weight parameter that can be used to give greater weight to more recent observations. In our simulations, we have used $\lambda = 0.99$, which gives a 74% weight to the observations over the 2001–2010 period, and a meaningful 22% to those over the 1991–2000 period.[7]

Having estimated all the relevant parameters, we draw a return for the first month (Jan/2011) for each of the six asset classes in our sample. Using the weights in the last three columns of Exhibit 1, we calculate the return for our three portfolios for that first month, and then we do the same for the subsequent 119 months. This yields one scenario, which consists of a series of 120 monthly returns for each portfolio over the Jan/2011-Dec/2020 period; Exhibit C1 shows one such scenario.

Finally, we repeat the whole process 10,000 times, thus generating 10,000 scenarios. Exhibits 3 and 4 summarize several aspects of the performance of S, G, and G_c portfolios over the 10,000 simulated horse races.

EXHIBIT C1

Simulation Methodology—One Scenario

This exhibit, one of the 10,000 scenarios of our simulations, shows the performance of three optimal portfolios, one selected by SRM (S), the other selected by GMM (G), and the other selected by GMM constrained to have weights no larger than 47.5% (G_C).

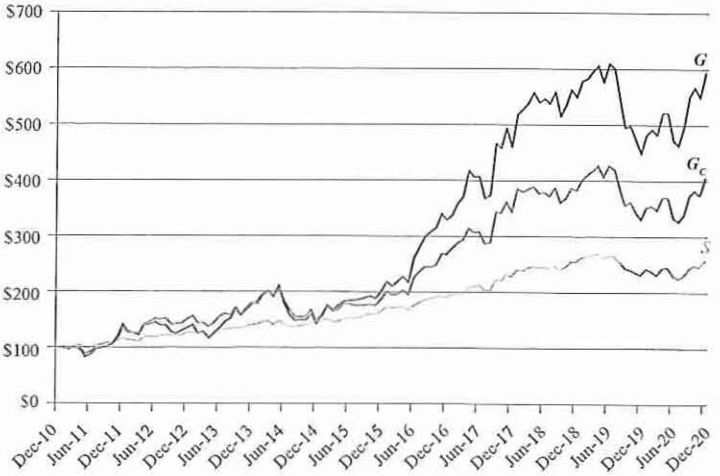

END NOTES

We would like to thank participants of the 7th Quantitative Asset & Risk Management Workshop (Venice, Italy) and the 2012 European FMA Conference (Istanbul, Turkey). Gabriela Giannattasio provided valuable research assistance. The views expressed below and any errors that may remain are entirely our own.

[1]As an example, consider two investments, one with a 5% certain return, and another with a 50-50 chance of a 200% gain or a 100% loss. Although this second alternative (with an expected value of 50%) may be, at least to some investors, more attractive than the first when making a one-time choice, it is a bad choice for *all* investors in a (long-term) multiperiod framework with reinvestment of gains and losses. This is the case because sooner or later, the 100% loss will occur and wipe out all the capital accumulated.

[2]However, Fama and MacBeth [1974] find substantial *economic* differences between the G portfolio and the market portfolio, the former having much higher (geometric mean) return and (beta) risk.

[3]We test for the equality in Sharpe ratios with the methodology proposed by Jobson and Korkie [1981] and Memmel [2003], and cannot reject the null hypothesis at the 5% level of significance.

[4]All the correlation coefficients we estimated are statistically significant, with the exception of two (between gold and U.S. stocks, and gold and U.S. real estate).

[5]Forecasting the parameters of the distributions of returns is one of the key technical issues. Once we have these distributions, by drawing a return every month, we get the monthly change of value for each asset class, and from these we compute the values of the three portfolios.

[6]See "RiskMetrics—Technical Document," fourth edition, 1996, chapter 5.

[7]In order to justify the 0.99 value note that with the frequently used $\lambda = 0.97$, the 1991–2000 and 2001–2010 periods would have had weights of 2% and 97%, respectively; with $\lambda = 0.95$, the same two periods would have had weights of 0.2% and 99.8%. We have explored the sensitivity of our results to changes in λ and found that they are not substantially affected.

REFERENCES

Adler, T., and M. Kritzman. "Mean-Variance versus Full-Scale Optimization: In and Out of Sample." *Journal of Asset Management*, 7 (2006), pp. 302-311.

Bekaert, G., C. Erb, C. Harvey, and T. Viskanta. "Distributional Characteristics of Emerging Market Returns and Asset Allocation." *The Journal of Portfolio Management*, Vol. 24, No. 2 (Winter 1998), pp. 102-116.

Bernstein, W., and D. Wilkinson. "Diversification, Rebalancing, and the Geometric Mean Frontier." Unpublished manuscript, 1997.

Christensen, M. "On the History of the Growth Optimal Portfolio." Unpublished manuscript, 2005.

Elton, E., and M. Gruber. "On the Maximization of the Geometric Mean with Lognormal Return Distribution." *Management Science*, Vol. 21, No. 4 (1974), pp. 483-488.

Estrada, J. "Mean-Semivariance Optimization: A Heuristic Approach." *Journal of Applied Finance* (Spring/Summer 2008), pp. 57-72.

——. "Geometric Mean Maximization: An Overlooked Portfolio Approach?" *The Journal of Investing*, Vol. 19, No. 4 (Winter 2010), pp. 134-147.

Fama, E., and J. MacBeth. "Long-Term Growth in a Short-Term Market." *The Journal of Finance*, Vol. 29, No. 3 (1974), pp. 857-885.

Grauer, R. "A Comparison of Growth Optimal and Mean Variance Investment Policies." *Journal of Financial and Quantitative Analysis*, Vol. 16, No. 1 (1981), pp. 1-21.

Hunt, B. "Growth Optimal Investment Strategy Efficacy: An Application on Long Run Australian Equity Data." *Investment Management and Financial Innovations*, 1 (2005), pp. 8-22.

Jobson, J.D., and B. Korkie. "Performance Hypothesis Testing with the Sharpe and Treynor Measures." *The Journal of Finance*, Vol. 36, No. 4 (1981), pp. 889-908.

Kelly, J. "A New Interpretation of Information Rate." *Bell System Technical Journal*, 35 (1956), pp. 917-926.

Kritzman, M., and D. Rich. "The Mismeasurement of Risk." *Financial Analysts Journal*, Vol. 58, No. 3 (May/June 2002), pp. 91-99.

Latane, H. "Criteria for Choice Among Risky Ventures." *Journal of Political Economy*, 67 (1959), pp. 144-155.

Lintner, J. "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." *Review of Economics and Statistics*, Vol. 47, No. 1 (1965), pp. 13-37.

Markowitz, H. "Portfolio Selection." *The Journal of Finance*, Vol. 7, No. 1 (1952), pp. 77-91.

———. *Portfolio Selection: Efficient Diversification of Investments.* New York: John Wiley & Sons, 1959.

———. "Investment for the Long Run: New Evidence for an Old Rule." *The Journal of Finance*, Vol. 31, No. 5 (1976), pp. 1273-1286.

Memmel, C. "Performance Hypothesis Testing with the Sharpe Ratio." *Finance Letters*, Vol. 1, No. 1 (2003), pp. 21-23.

Mossin, J. "Equilibrium in a Capital Asset Market." *Econometrica*, Vol. 34, No. 4 (1996), pp. 768-783.

Poundstone, W. *Fortune's Formula—The Untold Story of the Scientific Betting System That Beat the Casinos and Wall Street.* New York: Hill and Wang, 2005.

Roll, R. "Evidence on the 'Growth-Optimum' Model." *The Journal of Finance*, Vol. 28, No. 3 (1973), pp. 551-566.

Sharpe, W. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *The Journal of Finance*, Vol. 19, No. 3 (1964), pp. 425-442.

Treynor, J. "Toward a Theory of the Market Value of Risky Assets." Unpublished manuscript, 1961.

Weide, J., D. Peterson, and S. Maier. "A Strategy Which Maximizes the Geometric Mean Return on Portfolio Investments." *Management Science*, Vol. 23, No. 10 (1977), pp. 1117-1123.

Ziemba, W. "Note on Optimal Growth Portfolios When Yields Are Serially Correlated." *Journal of Financial and Quantitative Analysis*, Vol. 7, No. 4 (1972), pp. 1995-2000.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.