The Game-Theoretic Capital Asset Pricing Model

Vladimir Vovk

Glenn Shafer*

March 3, 2002

*Vovk is in the Department of Computer Science at Royal Holloway, University of London. Shafer is in the Graduate School of Management, Rutgers University.

The Game-Theoretic Capital Asset Pricing Model

ABSTRACT

Using Shafer and Vovk's game-theoretic framework for probability, we derive a capital asset pricing model from an efficient market hypothesis, with no assumptions about the beliefs or preferences of investors. Our efficient market hypothesis says that a speculator with limited means cannot beat a particular index by a substantial factor. The model we derive says that the difference between the average returns of a portfolio and the index should approximate the difference between the portfolio's covariance with the index and the index's variance. This leads to interesting new ways to evaluate the past performance of portfolios and funds.

The established general theory of capital asset pricing combines stochastic models for asset returns with a rich tapestry of economic ideas: no arbitrage, general equilibrium, and marginal utilities for current and future consumption (Campbell 2000, Cochrane 2001). Twenty years of work have demonstrated the power and flexibility of the combination; many different stochastic models and many different models for investors' marginal utility can be adopted, estimated, or predicted. There is little consensus, however, concerning the empirical validity of these different instantiations of the general theory. While this can be attributed in part to the very richness of the theory, which allows data on prices and returns to be looked at in different ways, it is also related to the theory's deep ambivalence about the meaning of its stochastic models. On the one hand, these models are hypotheses about the behavior of returns, to be compared with the empirical distribution of returns. On the other hand, they are hypotheses about investors' beliefs, which can be combined with hypotheses about investors' preferences in order to determine or predict asset prices. These two roles do not necessarily mesh. Investors can be mistaken about the future, and the fact that a stochastic model fits past asset returns does not go very far towards demonstrating that investors were using it to make their decisions.

One can try to achieve clarity within the established theory by making parsimonious assumptions about the stochastic process driving asset returns and about the marginal utility of investors. But this does not alleviate the problems arising from the multiple meanings of stochasticity. A more effective application of the principle of parsimony requires a deconstruction of stochasticity. One needs something more modest than the assumption that asset returns are generated by a stochastic process. In this article, we propose finding this something more modest in the game-theoretic framework recently advanced by Shafer and Vovk (2001). In this framework, limited opportunities to bet can be interpreted without any assumption of stochasticity. Shafer and Vovk show that the framework is adequate for the classical limit theorems of probability (the law of large numbers, the law of the iterated logarithm, and the central limit theorem), and that it can be used to make the theory and practice of option pricing more purely game-theoretic. In this article, we apply the framework to capital asset pricing.

In its simplest form, Shafer and Vovk's framework uses a two-player perfect-information sequential game. On each round, Player I can buy uncertain payoffs at given prices, and then Player II determines the values of the payoffs. The game, a precise and purely mathematical object, is connected to the world by an auxiliary nonmathematical hypothesis, Cournot's principle. Cournot's principle says that if Player I avoids risking bankruptcy, then he cannot multiply his initial capital in the game by a large factor. This principle gives empirical meaning to the game-theoretic forms of the classical limit theorems, for they say that certain approximations or convergences hold unless Player I is allowed to become very rich. The simplest form of the strong law of large numbers, for example, says that if there are infinitely many rounds of play, and Player I is allowed to make an arbitrary even-money bet on a binary outcome (heads or tails) on each round, then he has a strategy that does not risk bankruptcy and makes him infinitely rich if Player II does not make the proportion of heads converge to one-half. This theorem acquires empirical implications in any instance of coin tossing where we adopt Cournot's principle. Cournot's principle rules out Player I's becoming infinitely rich, and so we may conclude that Player II (reality) will make the proportion of heads converge to one-half.

A financial market provides a game of the required form: Player I is a speculator, who may buy various securities at set prices at the beginning of each trading period, and Player II is the market, which determines the securities' returns at the end of the period. If we measure Player I's capital relative to a particular market index, then Cournot's principle becomes an efficient market hypothesis: Player I cannot beat the index by a large factor. In this article, we show that this efficient market hypothesis implies an approximate relation between an investor's actual returns and the index's actual returns (Eq. (3) on p. 6) that resembles the equation for the security market line (an exact relation between theoretical quantities) in the classical Sharpe-Lintner capital asset pricing model (Sharpe 1964, Lintner 1965, Copeland and Weston 1988), the best-known instantiation of the established theory. Because of the resemblance, we call our model the *game-theoretic CAPM*.

Our efficient market hypothesis is, of course, consistent with the established theory. In its infinitary form, it can even be considered a consequence of the established theory. The established theory requires that the stochastic process for asset returns be absolutely continuous with respect to the risk-neutral probability measure obtained by normalizing state prices. Because asset prices are expected values for future returns with respect to the risk-neutral probability measure, this measure gives zero probability to the event that any given strategy for trading at these prices without risking bankruptcy will be infinitely successful. So the stochastic process must also give such an event zero probability. Thus our game-theoretic approach deconstructs rather than contradicts the established theory. It allows us to extract one part of the established theory—the efficient market hypothesis—and explore the consequences of this part alone.

While not contradicting the established theory, the game-theoretic CAPM differs from it radically in spirit. To avoid confusion, we need to keep three important aspects of the difference in view:

1. We make no assumptions whatsoever about the preferences or beliefs of investors.
2. We do not assume that asset returns are determined by a stochastic process. These returns are determined by the market, a player in our game. The market may act as it pleases, except that it is constrained in a certain sense by our efficient market hypothesis—our expectation that it will not allow spectacular success for any particular investment strategy that does not risk bankruptcy.
3. The predictions of our model concern the relation between the actual returns of an investor (or the actual returns of a security or portfolio) and the actual returns of an index. These predictions are precise enough to be confirmed or falsified by the actual returns, without any further modeling assumptions.

In this article, we check the predictions for several securities, and we find that they are usually correct.

The empirical success of our predictions, though modest, constitutes a challenge to the established theory. In spite of its parsimony, the game-theoretic CAPM can make reasonably precise and reasonably correct predictions concerning the relation between average return and empirical volatility and covariance. Can the established theory deliver enough more to give credibility to its much stronger assumptions? There is an analogy here with option pricing. Before Black-Scholes, it was customary to appeal to assumptions about investors' preferences and beliefs in order to derive prices for options. Now that these assumptions are seen as unnecessary, they are also seen as relatively dubious.

Our results can also be seen as a clarification of the roles of investors and speculators. An investor balances risk and return in an effort to balance present and future consumption, while a speculator is intent on beating the market. The established theory emphasizes the role of investors, but the efficient market hypothesis is usually justified by the presumed effectiveness of speculators. Speculators have already put so much effort into beating the market, the argument goes, that no opportunities remain for a new speculator who has no private information. The classical CAPM, still the most widely used instantiation of the established theory, bases its security market line, a relationship between expected return and covariance with the market, on the investor's effort to balance return with volatility, perceived as a measure of risk. Our game-theoretic CAPM, in contrast, shows that this relationship between return and covariance arises already from the speculator's elimination of opportunities to beat the market. So the relationship by itself does not provide any evidence that volatility measures risk, that it is perceived by the investor as doing so, or even that it can be predicted by the investor in advance.

In addition to providing an alternative understanding of the security market line, our results also lead to something entirely new: a new way of evaluating the past performance of portfolios and investors. According to our theory, the underperformance of a portfolio relative to the market index should be approximated by one-half the empirical variance of the difference between the return for the portfolio and the return for the index. We call this quantity the *theoretical performance deficit* (see Eq. (8) on p. 11). In the case of an investor or fund whose strategy cannot be sold short because it is not public information, the theoretical performance deficit should be a lower bound on the underperformance. Because a variance can be decomposed in many ways, the identification of the theoretical performance deficit opens the door to a plethora of new ways to analyze underperformance.

Because the game-theoretic apparatus in which our formal mathematical results are stated will be unfamiliar to most readers, and because these results include necessarily messy bounds on the errors in our approximations, we devote most of this article to informal statements and explanations. We state our results informally in Section I, and we explain the geometric intuition underlying them in Section II.

We present our formal theory in Sections III and IV. Section III introduces our game-theoretic framework, and Section IV presents our results as precise mathematical propositions within that framework. Section V illustrates how these propositions can be applied to data, and Section VI reviews the potential importance of our results.

Appendix A explains how the concepts of this article are related to the game-theoretic notion of probability developed by Shafer and Vovk (2001), and Appendix B provides proofs of the propositions stated in Section IV.

I. An Informal First Look

In this section we state the game-theoretic CAPM informally, say a few words about its derivation and its resemblance to the classical CAPM, and then explain how it leads to the theoretical performance deficit.

A. Average Return and Covariance

Consider a particular financial market and a particular market index m in which investors and speculators can trade. We assume that a speculator with limited means cannot beat the performance of m by a substantial factor; this is our *efficient market hypothesis* for m.

The *game-theoretic CAPM* for m, which follows from this hypothesis, says that if s is a security (or portfolio or other trading strategy) that can be sold short, then its average simple return, say μ_s, is approximated by

$$\mu_s \approx \mu_m - \sigma_m^2 + \sigma_{sm}, \tag{1}$$

where μ_m is the average simple return for the index m, σ_m^2 is the uncentered empirical variance of m's simple returns, and σ_{sm} is the uncentered empirical covariance of s's and m's simple returns. In order to make (1) into a mathematically precise statement, we must, of course, spell out just how close together μ_s and $\mu_m - \sigma_m^2 + \sigma_{sm}$ will be. We do this in Proposition 3 on p. 31.

If s cannot be sold short, then we obtain only

$$\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_{sm}. \tag{2}$$

This approximate inequality is made precise by Proposition 1 on p. 28. We call (1) the *long-short* game-theoretic CAPM, and we call (2) the *long* game-theoretic CAPM.

We can also write (1) in the form

$$\mu \approx (\mu_m - \sigma_m^2) + \sigma_m^2 \beta, \tag{3}$$

where we now write μ instead of μ_s for s's average return, and we write β for the ratio σ_{sm}/σ_m^2. We call the line $\mu = (\mu_m - \sigma_m^2) + \sigma_m^2 \beta$ in the (β, μ)-plane the *security market line* for the game-theoretic CAPM. We call β the *sensitivity* of s to m; it is the slope of the empirical regression through the origin of s's returns on m's returns.

B. The Empirical Nature of the Model

All the quantities in (1) are empirical: we are considering N trading periods, during which s has returns $s_1, \ldots, s_N$ and m has returns $m_1, \ldots, m_N$, and we have set

$$\begin{aligned} \mu_s &:= \frac{1}{N}\sum_{n=1}^{N} s_n, & \mu_m &:= \frac{1}{N}\sum_{n=1}^{N} m_n, \\ \sigma_m^2 &:= \frac{1}{N}\sum_{n=1}^{N} m_n^2, & \sigma_{sm} &:= \frac{1}{N}\sum_{n=1}^{N} s_n m_n. \end{aligned} \tag{4}$$

The s_n and m_n are simple returns; s_n is the total gain or loss (capital gain or loss plus dividends and redistributions) during period n from investing one monetary unit in s at the beginning of that period, and m_n is similarly the total gain or loss for m.

Our theory does not posit the existence of theoretical quantities that are estimated by the empirical quantities μ_s, μ_m, σ_m^2, and σ_{sm}, and there is nothing in our theory that requires these empirical quantities to be predictable in advance or stable over time.

Mathematical convenience in the development of our theory dictates that we use the uncentered definitions in (4) for σ_m^2 and σ_{sm}, so that β is the slope of the empirical linear regression through the origin. Numerically, however, we can expect (3) to remain valid if we use the centered counterparts of σ_m^2 and σ_{sm}, so that β is the slope of the usual empirical linear regression with a constant term, because there is usually little numerical difference between uncentered and centered empirical moments in the case of returns. The uncentered empirical variance σ_m^2 is related to its centered counterpart, $\frac{1}{N}\sum_n (m_n - \mu_m)^2$, by the identity

$$\frac{1}{N}\sum_n (m_n - \mu_m)^2 = \sigma_m^2 - \mu_m^2.$$

Because μ_m is usually of the same order of magnitude as σ_m^2 (see Section V), and because both are usually small, μ_m^2 will usually be much smaller and hence negligible compared to σ_m^2. Similarly,

$$\frac{1}{N}\sum_n (s_n - \mu_s)(m_n - \mu_m) = \sigma_{sm} - \mu_s\mu_m,$$

and $\mu_s\mu_m$ will also be negligible compared to σ_m^2. So a shift to the centered quantities will also make little difference in the ratio σ_{sm}/σ_m^2.

C. Why?

Proofs of Propositions 1 and 3 are provided in Appendix B, and the geometric intuition underlying them is explained in Section II. It may be helpful, however, to say a word here about the main idea.

Our starting point is the fact that the growth of an investment in s is best gauged not by its simple returns s_n but by its logarithmic returns $\ln(1+s_n)$ (see, e.g., Campbell, Lo, and MacKinlay 1997, p. 11). If we invest one unit in s at the beginning of the N periods, reinvest all dividends as we proceed, and write W_s for the resulting wealth at the end of N periods, then

$$\frac{1}{N}\ln W_s = \frac{1}{N}\ln\prod_{n=1}^{N}(1+s_n) = \frac{1}{N}\sum_{n=1}^{N}\ln(1+s_n).$$

So the Taylor expansion $\ln(1+x) \approx x - \frac{1}{2}x^2$ yields

$$\frac{1}{N}\ln W_s \approx \frac{1}{N}\sum_{n=1}^{N}\left(s_n - \frac{1}{2}s_n^2\right) = \mu_s - \frac{1}{2}\sigma_s^2. \tag{5}$$

We call $\frac{1}{N}\ln W \approx \mu - \frac{1}{2}\sigma^2$ the *fundamental approximation of asset pricing*. It shows us that investors and speculators should be concerned with volatility even if volatility does not measure risk, for volatility diminishes the final wealth that one might expect from a given average simple return. Moreover, it establishes approximate indifference curves in the (σ,μ)-plane for a speculator who is concerned only with final wealth. As we explain in Subsections II.B and II.D, we can reason about these indifference curves in much the same way as the classical CAPM reasons about an investor's mean-variance indifference curves (see, e.g., Copeland and Weston 1988, pp. 195–198), with similar results.

The imprecision of the approximations (1) and (2) arises partly from the imprecision of the fundamental approximation and partly from the imprecision of our efficient market hypothesis. We assume only that the market cannot be beat by a substantial factor, not that it cannot be beat at all.

D. Resemblance to the Classical CAPM

If we set

$$\mu_f := \mu_m - \sigma_m^2,$$

then we can rewrite (1) in the form

$$\mu_s \approx \mu_f + (\mu_m - \mu_f)\frac{\sigma_{sm}}{\sigma_m^2}. \tag{6}$$

This resembles the classical CAPM, which can be written as

$$\mathrm{E}(\tilde{R}_s) = R_f + (\mathrm{E}(\tilde{R}_m) - R_f)\frac{\mathrm{Cov}(\tilde{R}_s, \tilde{R}_m)}{\mathrm{Var}(\tilde{R}_m)}, \tag{7}$$

where R_f is the risk-free rate of return, and $\tilde{R}_s$ and $\tilde{R}_m$ are random variables whose realizations are the simple returns s_n and m_n, respectively (see Copeland and Weston 1988, Eq. (7.9) on p. 197). At a superficial level, we say that the game-theoretic CAPM modifies the classical CAPM in three ways:

1. It replaces theoretical expected values, variances, and covariances with empirical quantities. (The game-theoretic model has no probability measure and therefore no such theoretical quantities.)
2. It replaces an exact equation between theoretical quantities with an approximate equation between empirical quantities, with a precise error bound derived from the fundamental approximation and an efficient market hypothesis.
3. It replaces the risk-free rate of return with $\mu_m - \sigma_m^2$.

There are more fundamental differences, however. Because the left-hand side of the classical equation, Eq. (7), is the expected value of *s*'s future return, we might imagine an investor using this equation to determine a price for *s*. This justifies the name "capital asset pricing model" for the equation. In contrast, Eq. (6) is clearly not a model for the process by which capital assets are priced. It derives from a model for this process, the game described in Section III, together with an efficient market hypothesis. But in itself it is not a model for a process; it is merely a prediction about how empirical average simple returns and covariances will be related. It is an *ex post* rather than an *ex ante* model.

E. The Theoretical Performance Deficit

If we write W_m for the final wealth resulting from an initial investment of one unit in the index m and W_s for the final wealth of a particular investor who also begins with one unit capital, then

$$\frac{1}{N}\ln W_m - \frac{1}{N}\ln W_s \overset{\text{FA}}{\approx} \left(\mu_m - \frac{1}{2}\sigma_m^2\right) - \left(\mu_s - \frac{1}{2}\sigma_s^2\right) \overset{\text{CAPM}}{\approx} \frac{1}{2}\sigma_s^2 - \sigma_{sm} + \frac{1}{2}\sigma_m^2 = \frac{1}{2}\sigma_{s-m}^2.$$

Here FA indicates use of the fundamental approximation, $\frac{1}{N}\ln W \approx \mu - \frac{1}{2}\sigma^2$, and CAPM indicates use of the game-theoretic CAPM, $\mu_s - \mu_m \approx \sigma_{sm} - \sigma_m^2$. The final step uses the identity $\sigma_{s-m}^2 = \sigma_s^2 - 2\sigma_{sm} + \sigma_m^2$, where $s - m$ is the vector of differences in the returns: $s - m = (s_1 - m_1, \ldots, s_N - m_N)$.

So when an investor holds a fixed portfolio or follows some other strategy that can be sold short, we should expect

$$\frac{1}{N}\ln W_m - \frac{1}{N}\ln W_s \approx \frac{1}{2}\sigma^2_{s-m}, \tag{8}$$

and even when s cannot be sold short, we should expect

$$\frac{1}{N}\ln W_m - \frac{1}{N}\ln W_s \gtrapprox \frac{1}{2}\sigma^2_{s-m}. \tag{9}$$

In words: s's average logarithmic return can be expected to fall short of m's by approximately $\sigma^2_{s-m}/2$, or by even more if there are difficulties in short selling. The approximation (8) is made precise by Proposition 4 on p. 31, and the approximate inequality (9) is made precise by Proposition 2 on p. 30.

We call $\sigma^2_{s-m}/2$ the *theoretical performance deficit* for s. If we consider the market index m a maximally diversified portfolio, then s's theoretical performance deficit can be attributed to insufficient diversification.

It is natural to decompose the vector of simple returns s into a part in the direction of the vector m and a part orthogonal to m: $s = \beta m + e$. Then we have $s - m = (\beta - 1)m + e$, and

$$\sigma^2_{s-m} = (\beta - 1)^2\sigma^2_m + \sigma^2_e.$$

Thus s's theoretical performance deficit, $\sigma^2_{s-m}/2$, decomposes into two parts:

$$\text{deficit due to nonunit sensitivity to } m\text{: } \frac{1}{2}(\beta - 1)^2\sigma^2_m, \tag{10}$$

and

$$\text{deficit due to volatility orthogonal to } m\text{: } \frac{1}{2}\sigma^2_e. \tag{11}$$

These two parts of the deficit represent two aspects of insufficient diversification. Many other decompositions of $\sigma^2_{s-m}/2$ are possible, corresponding to events inside and outside the market.

Such decompositions may be useful for analyzing and comparing the performance of different mutual funds, especially funds that do try to track the market.

There is nothing in our theory that would require the theoretical performance deficit of a particular security or portfolio to persist from one period of time to another. On the contrary, a persistence that is too predictable and substantial would give a speculator an opportunity to beat the market by shorting that security or portfolio, thus contradicting our efficient market hypothesis.

In the case of an investor or fund whose strategy cannot be shorted because it is not public information, persistence of the theoretical performance deficit or certain components of that deficit cannot be ruled out. It would be interesting to study the extent to which such persistence occurs.

II. The Geometric Intuition

In this section, we explain the geometric intuition that underlies the game-theoretic CAPM. This explanation will be repeated in a terser and more formal way in the proofs in Appendix B.

Here is a summary. We begin with what we call the *capital market parabola* for m: the curve in the (σ, μ)-plane consisting of all volatility-return pairs that yield approximately the same final wealth as m. The efficient market hypothesis for m says that the volatility-return pair for the simple returns s achieved by any given investor should fall under the capital market parabola for m, as should the volatility-return pair for any particular mixture of m and s. In order for this to be true for mixtures that contain mostly m and only a little s, the trajectory traced by the volatility-return pair as s's share in the mixture approaches zero must be approximately tangent to the parabola. The formula that expresses this conclusion turns out to be our CAPM: $\mu_s \approx \mu_m - \sigma_m^2 + \sigma_{sm}^2$. The conclusion requires that short selling of s be possible, so that the mixture can include a negative amount of s; otherwise we can conclude only that the trajectory cannot approach the parabola from above, and this yields only $\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_{sm}^2$.

We should not lose sight of the approximate and heuristic nature of this argument. There are two sources of inexactness. First, the capital market parabola is only approximately an indifference curve for total wealth; this is the fundamental approximation. Second, the efficient market hypothesis for m is itself only approximately correct. In this section, we ignore these two sources of inexactness. In Sections III and IV we analyze them carefully, so as to replace the vague approximations of this and the preceding section with inequalities that involve precise error bounds.

A. The Capital Market Parabola

As we saw in Subsection I.C, a speculator who is concerned only with his final wealth will be roughly indifferent between volatility-return pairs that have the same value of $\mu - \frac{1}{2}\sigma^2$—i.e., volatility-return pairs that lie on the same parabola $\mu = \frac{1}{2}\sigma^2 + c$. Figure 1 depicts two parabolas of this form in the half-plane consisting of (σ, μ) with $\sigma > 0$. The parabola that lies higher in the figure corresponds to a higher level of final wealth.

Insert Figure 1 about here.

The efficient market hypothesis for the market index m implies that the volatility-return pair achieved by a particular investor should lie approximately on or below the final wealth parabola on which (σ_m, μ_m) lies. This is the parabola

$$\mu = \frac{1}{2}\sigma^2 + \left(\mu_m - \frac{1}{2}\sigma_m^2\right),$$

the *capital market parabola* (CMP) for m.

In general, the parabola that goes through the volatility-return pair for a particular security or portfolio s,

$$\mu = \frac{1}{2}\sigma^2 + \left(\mu_s - \frac{1}{2}\sigma_s^2\right), \tag{12}$$

intersects the μ-axis at $\mu_s - \frac{1}{2}\sigma_s^2$. Because this is the constant simple return that gives approximately the same final wealth as s, we call it s's *volatility-free equivalent*.

Strictly speaking, a constant simple return μ does not have zero volatility when we use the uncentered definition; its volatility is

$$\sigma := \sqrt{\frac{1}{N}\sum_{n=1}^{N} \mu^2} = |\mu|.$$

This is why the indifference curves in Figure 1 do not quite reach the μ-axis; they stop at the line $\mu = \sigma$ above the σ-axis and at the line $\mu = -\sigma$ below the σ-axis. But the height of parabola (12)'s intersection with this line will be practically the same as the height of its intersection with the μ-axis.

B. Mixing *s* and *m*: The Long CAPM

We are now in a position to derive the approximate inequality $\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_{sm}$ from our efficient market hypothesis. To do this, it suffices to suppose that a speculator is allowed to hold long positions in s and m. We do not need to suppose that he can also sell s short.

Insert Figure 2 about here.

Suppose the speculator maintains a portfolio p that mixes s and m, say ε of s and $(1-\varepsilon)$ of m, where $0 \leq \varepsilon \leq 1$. (He rebalances at the beginning of every period so that s always accounts for the fraction ε of p's capital.) Under our efficient market hypothesis, the volatility-return pair for p lies approximately on or below the CMP no matter what the value of ε is. As ε varies between 0 and 1, (σ_p, μ_p) traces a trajectory, perhaps as indicated in Figure 2. We will usually think of this trajectory as running in the direction from $\varepsilon = 1$ down to $\varepsilon = 0$—i.e., from (σ_s, μ_s) somewhere below the CMP to (σ_m, μ_m) on the CMP. We have

$$\mu_p = \varepsilon\mu_s + (1-\varepsilon)\mu_m \tag{13}$$

and

$$\sigma_p = \sqrt{\varepsilon^2\sigma_s^2 + 2\varepsilon(1-\varepsilon)\sigma_{sm} + (1-\varepsilon)^2\sigma_m^2}. \tag{14}$$

Hence

$$\frac{\partial\mu_p}{\partial\varepsilon}\bigg|_{\varepsilon=0} = \mu_s - \mu_m$$

and

$$\frac{\partial\sigma_p}{\partial\varepsilon}\bigg|_{\varepsilon=0} = \frac{\sigma_{sm} - \sigma_m^2}{\sigma_m}.$$

(Cf. Copeland and Weston 1988, p. 197.) If the second of these two derivatives is nonzero, then their ratio,

$$\frac{\mu_s - \mu_m}{(\sigma_{sm} - \sigma_m^2)/\sigma_m}, \tag{15}$$

is the slope of the tangent to the trajectory at (σ_m, μ_m).

Our goal is to understand why the long CAPM should hold—i.e., to understand why

$$\mu_s - \mu_m \leq \sigma_{sm} - \sigma_m^2 \tag{16}$$

should hold approximately. To this end, we consider four cases:

1. $\mu_s - \mu_m \leq 0$ and $\sigma_{sm} - \sigma_m^2 \geq 0$.
2. $\mu_s - \mu_m \geq 0$ and $\sigma_{sm} - \sigma_m^2 \leq 0$, but not both are equal to 0.
3. $\mu_s - \mu_m > 0$ and $\sigma_{sm} - \sigma_m^2 > 0$.
4. $\mu_s - \mu_m < 0$ and $\sigma_{sm} - \sigma_m^2 < 0$.

Any two real numbers are related to each other in one of these four ways.

In Case 1, we obtain (16) immediately: a nonpositive quantity cannot exceed a nonnegative one. Figure 2 is an example of this case. We see from the figure that μ_s is below μ_m, and that the trajectory approaches (σ_m, μ_m) from the southeast. So $\mu_s - \mu_m$ is strictly negative and the slope (15) is negative; it follows that $\sigma_{sm} - \sigma_m^2$ is positive.

Case 2 is ruled out by the efficient market hypothesis for m. It tells us that μ_s is at least as large as μ_m, and because μ_p changes monotonically with ε, this means that the trajectory must approach (σ_m, μ_m) from above or the side. It also tells us that the slope (15) is negative unless one of the quantities is zero. So the trajectory approaches (σ_m, μ_m) from the northwest (directly from the west if $\mu_s - \mu_m = 0$, directly from the north if $\sigma_{sm} - \sigma_m^2 = 0$). This means approaching (σ_m, μ_m) from above the CMP, in contradiction to our efficient market hypothesis.

In Case 3, the slope of the trajectory at (σ_m, μ_m) is positive, and the trajectory approaches (σ_m, μ_m) from the northeast. Because the trajectory must lie under the CMP, its slope at (σ_m, μ_m) cannot exceed the CMP's slope at (σ_m, μ_m), which is σ_m:

$$\frac{\mu_s - \mu_m}{(\sigma_{sm} - \sigma_m^2)/\sigma_m} \leq \sigma_m.$$

Multiplying both sides by the denominator, we obtain (16).

Case 4 is similar to Case 3; the slope is again positive, but now the approach is from the southwest, and so staying under the CMP requires that the slope be at least as great:

$$\frac{\mu_s - \mu_m}{(\sigma_{sm} - \sigma_m^2)/\sigma_m} \geq \sigma_m.$$

This time the denominator is negative, and so multiplying both sides by it again yields (16).

C. The Capital Market Line

We should pause to note that the approximate inequality that we have just argued for,

$$\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_{sm}, \tag{17}$$

implies a strengthening of the statement that (σ_s, μ_s) should be approximately on or below the capital market parabola in the (σ, μ)-plane. This pair should also be approximately on or below the line tangent to this parabola at (σ_m, μ_m). (See Figure 3.)

Insert Figure 3 about here.

To see this, it suffices to rewrite (17) in the form

$$\mu_s \lessapprox \mu_m - \sigma_m^2 + \rho_{sm}\sigma_m\sigma_s,$$

where ρ_{sm} is the uncentered correlation coefficient between s and m. Because $\rho_{sm} \leq 1$, this implies

$$\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_m\sigma_s. \tag{18}$$

In other words, (σ_s, μ_s) must lie approximately on or below the line

$$\mu = (\mu_m - \sigma_m^2) + \sigma_m\sigma. \tag{19}$$

This line, which we call the *capital market line* (CML), is the tangent to the CMP at (σ_m, μ_m).

D. Shorting s to Go Longer in m: The Long-Short CAPM

The trajectories we have been considering trace what happens for different values of the fraction ε occupied on a portfolio p by s. So far, we have assumed that $0 \leq \varepsilon \leq 1$. But if our speculator is allowed to short s in order to go longer in m, then he can take ε past zero into negative territory. This means extending the trajectory in the direction it is pointing as it approaches (σ_m, μ_m).

Insert Figure 4 about here.

We evidently have a problem if the trajectory approaches the CMP as in Figure 2. In such a case, extending the trajectory past (σ_m, μ_m) by going short in s a small amount ε means extending the trajectory above the CMP, in contradiction to our efficient market hypothesis. So such trajectories are ruled out when the speculator is allowed to sell s short.

There are only two conditions under which selling s short by a small amount ε will not move the speculator above the CMP:

1. If the partial derivatives (13) and (14) are both zero, then selling s short by a small amount ε will have no first-order effect; the pair (σ_p, μ_p) will remain approximately equal to (σ_m, μ_m).
2. If the trajectory is approximately tangent to the CMP at (σ_m, μ_m), as in Figure 4, then the speculator will remain under the CMP even if he can extend the trajectory a small amount past (σ_m, μ_m).

The long-short CAPM,

$$\mu_s \approx \mu_m - \sigma_m^2 + \sigma_{sm},$$

holds under both conditions. It holds under the first condition because $\mu_s - \mu_m$ and $\sigma_{sm} - \sigma_m^2$ are both zero. It holds under the second condition because the slope (15) is approximately σ_m.

It may be helpful to elaborate some further implications of the first of the two conditions. From $\mu_s - \mu_m = 0$, we find that μ_p is constant: $\mu_p = \mu_s = \mu_m$. From $\sigma_{sm} - \sigma_m^2 = 0$, we find that $s = m + e$, where e is orthogonal to m, so that $p = m + \varepsilon e$ and $\sigma_p^2 = \sigma_m^2 + \varepsilon^2 \sigma_e^2$. Geometrically, this means that the trajectory approaches (σ_m, μ_m) directly from the east as ε moves from 1 down to 0, and then eventually moves directly back east as ε moves substantially into negative territory.

III. Quantifying Our Efficient Market Hypotheses

The efficient market hypotheses used in this article assert that a speculator with limited means cannot beat a specified market index m by a substantial factor. He cannot achieve a final wealth many times greater than what is achieved by someone who invests the same limited initial capital in m.

At first glance, we might doubt whether such a hypothesis could stand up to the facts. No matter what market, what period of time, and what index m we choose, we can retrospectively find strategies and perhaps even securities that do beat m by a substantial factor. A strategy that shifts at the beginning of each day to those securities that increase in price the most that day will usually beat any index spectacularly. So what do we mean when we say that a speculator cannot beat m by a substantial factor? We mean that we do not expect any particular speculator (or any particular security, portfolio, or strategy selected in advance) to do much better than the market. We do not expect the speculator's final wealth to exceed by a large factor the final wealth that he would have achieved simply by investing his initial wealth in the market index m. The larger the factor, the stronger our expectation. If α is a positive number very close to zero, and the speculator starts with initial wealth equal to one monetary unit, then we strongly expect his final wealth will be less than $\frac{1}{\alpha}W_m$, where W_m is the final wealth obtained by investing one monetary unit in m at the outset. This is an expectation about the market's behavior: the market will follow a course that makes the speculator's wealth less than $\frac{1}{\alpha}W_m$.

In this section, we review some ideas from Shafer and Vovk (2001), where this way of quantifying efficient market hypotheses is given a natural game-theoretic foundation. In Subsection A, we formulate the *basic capital asset pricing game (basic CAPG)*. In Subsections B and C, we discuss how this game, in itself only a mathematical object, can be used to model securities markets. In Subsection D, we define two variations on the basic CAPG, which provide the settings for the precise mathematical formulations of the long CAPM and the long-short CAPM that we present later, in Section IV.

A. The Basic Capital Asset Pricing Game

The capital asset pricing game has two principal players, SPECULATOR and MARKET, who alternate play. On each round, SPECULATOR decides how much of each security in the market to hold (and possibly short), and then MARKET determines SPECULATOR's gain by deciding how the prices of the securities change. Allied with MARKET is a third player, INVESTOR, who also invests each day. The game is a perfect-information game: each player sees the others' moves.

We assume that there are $K+1$ securities in the market and N rounds (trading periods) in the game. We number the securities from 0 to K and the rounds from 1 to N, and we write x_n^k for the simple return on security k in round n. For simplicity, we asume that $-1 < x_n^k < \infty$ for for all k and n; a security price never becomes zero. We write x_n for the vector $(x_n^0, \ldots, x_n^K)$, which lies in $(-1,\infty)^{K+1}$. MARKET determines the returns; x_n is his move on the nth round.

We assume that the first security, indexed by 0, is our market index m; thus x_n^0 is the same as m_n, the simple return of the market index m in round n. If m is a portfolio formed from the other securities, then x_n^0 is an average of the $x_n^1, \ldots, x_n^K$, but we do not insist on this.

We write $\mathcal{M}_n$ for the wealth at the end of round n resulting from investing one monetary unit in m at the beginning of the game:

$$\mathcal{M}_n := \prod_{i=1}^{n}(1+x_i^0) = \prod_{i=1}^{n}(1+m_i).$$

Thus $\mathcal{M}_N$ is the final wealth resulting from this investment. This is the quantity we earlier designated by W_m.

INVESTOR begins with capital equal to one monetary unit and is allowed to redistribute his current capital across all $K+1$ securities on each round. If we write $\mathcal{G}_n$ for his wealth at the end of the nth round, then

$$\mathcal{G}_n := \prod_{i=1}^{n} \sum_{k=0}^{K} g_i^k (1 + x_i^k),$$

where g_i^k is the fraction of his wealth he holds in security k during the ith round. The g_i^k must sum to 1 over k, but g_i^k may be negative for a particular k (in this case INVESTOR is selling k short). INVESTOR's final wealth is $\mathcal{G}_N$. Thus $\mathcal{G}_N$ is the same as what we earlier called W_s. We will also write s_n for INVESTOR's simple return on round n:

$$s_n := \frac{\mathcal{G}_n - \mathcal{G}_{n-1}}{\mathcal{G}_{n-1}} = \sum_k g_n^k x_n^k. \tag{20}$$

We call the set of all possible sequences $(g_1, x_1, \ldots, g_N, x_N)$ the *sample space* of the game, and we designate it by Ω:

$$\Omega := \left(\mathbb{R}^{K+1} \times (-1, \infty)^{K+1}\right)^N.$$

We call any subset of Ω an *event*. Any statement about INVESTOR's returns determines an event, as does any comparison of INVESTOR's and MARKET's returns.

SPECULATOR also starts with one monetary unit and is allowed to redistribute his current capital across all $K+1$ securities on each round. We write $\mathcal{H}_n$ for his wealth at the end of the nth round:

$$\mathcal{H}_n := \prod_{i=1}^{n} \sum_{k=0}^{K} h_i^k (1 + x_i^k),$$

where h_i^k is the fraction of his wealth he holds in security k during the ith round. The moves by SPECULATOR are not recorded in the sample space; they do not define events.

To complete the specification of the game, we select a number α and an event A, and we agree that SPECULATOR will win the game if he beats the index by the factor $\frac{1}{\alpha}$ or if A happens. The number α is our *significance level*, and the event A is SPECULATOR's *auxiliary* *goal*. This auxiliary goal might, for example, be the event that INVESTOR's average simple return μ_s approximates $\mu_m - \sigma_m^2 + \sigma_{sm}$ to some specified accuracy.

Putting all these elements together, we have this game:

BASIC CAPITAL ASSET PRICING GAME (BASIC CAPG)

Players: INVESTOR, MARKET, SPECULATOR

Parameters:

Natural number K (number of non-index securities in the market)

Natural number N (number of rounds or trading periods)

Real number α satisfying $0 < \alpha \leq 1$ (significance level)

$A \subseteq \Omega$ (auxiliary goal)

Protocol:

$\mathcal{G}_0 := 1$.

$\mathcal{H}_0 := 1$.

$\mathcal{M}_0 := 1$.

FOR $n = 1, 2, \ldots, N$:

INVESTOR selects $g_n \in \mathbb{R}^{K+1}$ such that $\sum_{k=0}^{K} g_n^k = 1$.

SPECULATOR selects $h_n \in \mathbb{R}^{K+1}$ such that $\sum_{k=0}^{K} h_n^k = 1$.

MARKET selects $x_n \in (-1, \infty)^{K+1}$.

$\mathcal{G}_n := \mathcal{G}_{n-1} \sum_{k=0}^{K} g_n^k (1 + x_n^k)$.

$\mathcal{H}_n := \mathcal{H}_{n-1} \sum_{k=0}^{K} h_n^k (1 + x_n^k)$.

$\mathcal{M}_n := \mathcal{M}_{n-1} (1 + x_n^0)$.

Winner: SPECULATOR wins if $\mathcal{H}_n \geq 0$ for $n = 1, \ldots, N$ and either (1) $\mathcal{H}_N \geq \frac{1}{\alpha} \mathcal{M}_N$ or (2) $(g_1, x_1, \ldots, g_N, x_N) \in A$. Otherwise INVESTOR and MARKET win.

The requirement that SPECULATOR keep $\mathcal{H}_n$ nonnegative in order to win formalizes the idea that he has limited means. It ensures that when $\mathcal{H}_N \geq \frac{1}{\alpha}\mathcal{M}_N$, he really has turned an initial capital of only one monetary unit into $\frac{1}{\alpha}\mathcal{M}_N$. If he were allowed to continue on to the $(n+1)$st round when $\mathcal{H}_n < 0$, he would be borrowing money—i.e., drawing on a larger capital—and if he then finally achieved $\mathcal{H}_N \geq \frac{1}{\alpha}\mathcal{M}_N$, it would not be fair to credit him with doing so with his limited initial means of only one monetary unit. Because SPECULATOR must keep $\mathcal{H}_n$ always nonnegative in order to win, a strategy for SPECULATOR cannot guarantee his winning if it permits the other players to force $\mathcal{H}_n < 0$ for some n. In other words, a winning strategy for SPECULATOR cannot risk bankruptcy.

Formally, the basic CAPG allows SPECULATOR to sell securities short. However, if SPECULATOR sells security k short on round n, then MARKET has the option of making the return x_n^k so large that $\mathcal{H}_n$ becomes negative, resulting in SPECULATOR's immediately losing the game. So no winning strategy for SPECULATOR can involve short selling. In Subsection D, we discuss how the rules of the game can be modified to make short selling a real possibility for SPECULATOR.

B. Predictions from the Efficient Market Hypothesis

In order for SPECULATOR to win our game, *either* he must become very rich relative to the market index m (he beats m by the factor $\frac{1}{\alpha}$) or *else* the event A must happen. In the next section, we will show that for certain choices of A and α, SPECULATOR can win—he has a winning strategy. But our efficient market hypothesis predicts that the market will not allow him to become very rich relative to m, and this implies that A will happen. In this sense, our efficient market hypothesis predicts that A will happen.

To formalize this idea, we make the following definition: The efficient market hypothesis for *m predicts* the event *A at level* α if Speculator has a winning strategy in the basic CAPG with A as the auxiliary goal and α as the significance level.

As we explained earlier, our confidence that SPECULATOR will not beat the market by $\frac{1}{\alpha}$ is greater for smaller α. So a prediction of A at level α becomes more emphatic as α decreases.

C. Is the Game Realistic?

It may be unclear to the reader how this game can be used as a model. Is a securities market really a perfect-information game? Does it involve only three players?

We relate the game to an actual securities market by thinking of INVESTOR as a particular individual investor or fund. INVESTOR may do whatever a real investor may do: he may follow some particular static strategy (hold only a particular security or portfolio); he may follow some particular dynamic strategy; or he may play opportunistically, without any strategy chosen in advance. MARKET represents all the other participants in the market. Because MARKET and INVESTOR play the game as a team against SPECULATOR, we can even think of MARKET as representing all the participants in the market, including INVESTOR.

SPECULATOR does not represent a real investor. Rather, he represents the hypothetical or imaginary investor referred to by our efficient market hypothesis. Our efficient market hypothesis says that a speculator cannot multiply his initial capital by a substantial factor relative to the index m. Conceptually, this is not necessarily a statement about a real investor. It is a statement about strategies: we do not expect any particular strategy selected in advance to beat m by a substantial factor.

The roles of INVESTOR and SPECULATOR should be intuitively clear from Section II. INVESTOR earns the simple returns s. SPECULATOR forms a portfolio p by mixing ε of s and $1-\varepsilon$ of m. We have SPECULATOR move after INVESTOR so that he knows what INVESTOR is doing with his capital and can replicate it with ε of his own capital. Of course, there are circumstances in which SPECULATOR can replicate what INVESTOR is doing without knowing exactly what it is. For example, INVESTOR might represent a fund in which SPECULATOR can invest ε of his capital. Our results will also apply to this case.

The winning strategies for SPECULATOR that we construct to prove our propositions are all of the simple form that we just mentioned: SPECULATOR mixes INVESTOR's moves with *m*, perhaps going short in INVESTOR's moves to go longer in *m*. Because these simple strategies are sufficient, the efficient market hypothesis that we need in order to draw our practical conclusions from the propositions is sometimes relatively weak. Instead of assuming that no speculator can beat the market by a large factor, no matter how smart and imaginative he is, it is enough to assume that no speculator can beat *m* by a large factor using strategies at most slightly more complicated than those used by the investors or funds whose performance we are studying.

We have said that MARKET represents all the actual agents in the market aside, perhaps, from INVESTOR. The reader might protest that these agents are unlikely to function as a single player capable of acting strategically. Our assumption of perfect information may also be puzzling. Most of the agents who constitute MARKET will not observe INVESTOR's moves. What does it mean to say that they observe moves by the imaginary player SPECULATOR? What does it mean to say that all the agents, INVESTOR together with the agents constituting MARKET, are playing against this imaginary player?

These puzzles disappear once the form of our mathematical results is understood. All our results say that the efficient market hypothesis for *m* predicts a particular event *A* at a particular level α. (See, for example, Proposition 1 on p. 28.) This means that SPECULATOR has a winning strategy in our game for certain values of the parameters. Such results obviously remain valid if we modify the game by restricting the knowledge and freedom of action of INVESTOR and MARKET; if SPECULATOR can win the game as described, he can also win any game in which his opponents are weaker. Because we endow MARKET and INVESTOR with so much knowledge and freedom of action in the game as described, our results say that SPECULATOR can achieve his goal no matter how strategically MARKET plays, no matter how much INVESTOR and MARKET know, and no matter how much INVESTOR and MARKET collude.

D. The Long and Long-Short Capital Asset Pricing Games

We do not actually use the basic CAPG for our mathematical work in the next section. Instead, we use two variations, which we call the *long CAPG* and the *long-short CAPG*.

Both the long CAPG and the long-short CAPG are obtained from the basic CAPG by restricting how the players can move:

- The *long CAPG* is obtained by replacing the conditions $g_n \in \mathbb{R}^{K+1}$ and $h_n \in \mathbb{R}^{K+1}$ in the protocol for the basic CAPG by the conditions $g_n \in [0,\infty)^{K+1}$ and $h_n \in [0,\infty)^{K+1}$, respectively. In other words, both INVESTOR and SPECULATOR are forbidden to sell securities short.
- The *long-short CAPG* has two extra parameters: a positive constant C (perhaps very large), and a positive constant δ (perhaps very small). It is obtained by replacing the condition $g_n \in \mathbb{R}^{K+1}$ in the protocol for the basic CAPG by the condition $g_n \in [0,\infty)^{K+1}$ and replacing the condition $x_n \in (-1,\infty)^{K+1}$ by the conditions $x_n \in (-1,C]^{K+1}$ and $m_n \geq -1+\delta$. (Remember that $m_n = x_n^0$.) In other words, INVESTOR is not allowed to sell short, and MARKET is constrained so that an individual security cannot increase too much in value on a single round and the market index m cannot lose too much of its value on a single round. These constraints on INVESTOR and MARKET make it possible for SPECULATOR to go short in INVESTOR's moves, at least a bit, without risking bankruptcy.

The concept of prediction is defined for these games just as for the basic CAPG: The efficient market hypothesis for *m predicts A at level* α for one of the games if SPECULATOR has a winning strategy in that game with A as the auxiliary goal and α as the significance level. In Subsection IV.A we show that certain events are predicted at level α in the long CAPG, and in Subsection IV.C we show that certain events are predicted at level α in the long-short CAPG.

Our results for the long CAPG in Subsection IV.A do not actually require forbiding SPECULATOR's selling short; it is enough to forbid INVESTOR's selling short. These results say that

SPECULATOR has a winning strategy for certain values of the parameters, and hence they cannot be affected by whether we formally allow SPECULATOR to sell short. Because MARKET remains unconstrained in the long CAPG, the lesson we learned for the basic CAPG on p. 23 applies: No winning strategy for SPECULATOR can go short, because MARKET can bankrupt him whenever he does go short.

We can also weaken the constraint that INVESTOR not sell short in the long CAPG. It is enough to constrain INVESTOR and MARKET to choose g_n and x_n so that $s_n > -1$ (see Eq. (20) on p. 21). This ensures that they cannot bankrupt SPECULATOR if he takes only long positions in s and m.

We can similarly weaken the constraints on INVESTOR and MARKET in the long-short CAPG: Require only that (1) $m_n \geq -1+\delta$ (the market index never drops too much on a single round) and (2) $-1 < s_n \leq C$ (INVESTOR never becomes bankrupt and never makes too great a return on a single round).

IV. Precise Mathematical Results

We now state propositions that express precisely, within the game-theoretic framework, the assertions that we outlined informally in Section I. Proofs of these propositions are provided in Appendix B.

A. The Long CAPM

Our first proposition translates the approximate inequality that we call the long CAPM, Eq. (2), into a precise inequality.

Proposition 1 *For any* $\alpha \in (0,1]$ *and any* $\varepsilon \in (0,1]$*, the efficient market hypothesis for m predicts*

$$\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm} < \frac{E}{\varepsilon} + \frac{\ln \frac{1}{\alpha}}{N\varepsilon} + \frac{\varepsilon \sigma_{s-m}^2}{2} \tag{21}$$

at level α *in the long CAPG, where*

$$E := \frac{1}{N} \sum_{n=1}^{N} \Big(\Gamma(m_n) - \gamma((1-\varepsilon)m_n + \varepsilon s_n) \Big) \tag{22}$$

and the functions Γ *and* γ *are defined by*

$$\Gamma(x) := \frac{1}{3}x^3, \qquad \gamma(x) := \frac{1}{3}\left(\frac{x}{1+x}\right)^3.$$

The quantity E bounds the accuracy of the fundamental approximation. It is awkwardly complicated because we have made the bound as tight as possible. In theory, E can be negative, but it is typically positive, and certainly the right-hand side of (21) as a whole is typically positive.

Although Proposition 1 is valid as stated, for any natural number N, any $\alpha \in (0,1]$, and any $\varepsilon \in (0,1]$, its theoretical significance is greatest when these parameters are chosen so that the right-hand side of (21) is small in absolute value relative to the typical size of the individual terms on the left-hand side, μ_s, μ_m, σ_m^2, and σ_{sm}. When this is so, (21) can be read roughly as $\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm} \lessapprox 0$, or $\mu_s \lessapprox \mu_m - \sigma_m^2 + \sigma_{sm}$.

In this paragraph, we will use the phrase *relatively small* to mean "small in absolute value relative to the typical size of μ_s, μ_m, σ_m^2, and σ_{sm}". In order for the right-hand side of (21) to be relatively small, we need all three of its terms to be relatively small. To see what this involves, let us look at these three terms individually:

- The theoretical performance deficit $\sigma_{s-m}^2/2$, which measures s's lack of diversification, is typically of the same order of magnitude as μ_s, μ_m, σ_m^2, and σ_{sm}. So we need to make ε small.

- To make our efficient market hypothesis realistic, we must choose α significantly less than one. So in order to make the term $\frac{\ln(1/\alpha)}{N\epsilon}$ relatively small, we must make the number of rounds N large even relative to $1/\epsilon$. Because the typical size of μ_s, μ_m, σ_m^2, and σ_{sm} decrease when the time period for each round is made shorter, it is not enough to make N large by making these individual time periods short. We must make the total period of time studied long.

- Once we have chosen a small ϵ, we must make E extremely small in order to make E/ϵ relatively small. Because E is essentially the difference between two averages of the third moments of the returns, we can make it extremely small by making the individual trading periods sufficiently short.

To summarize, we can hope to get a tight bound in (21) only if we choose ϵ small and consider frequent returns (perhaps daily returns) over a long period of time.

These points can be made much more clearly by a more formal analysis of the asymptotics. Fix arbitrarily small $\alpha > 0$ and $\epsilon > 0$. (We make ϵ small because we need it small; we make α small to show that we can tolerate it small.) Suppose trading happens during an interval of time $[0,T]$ that is split into N subintervals of length $dt = T/N$, and let $T \to \infty$ and $dt \to 0$. We can expect that s_n and m_n will have the order of magnitude $(dt)^{1/2}$, E will have the order of magnitude $(dt)^{3/2}$, and μ_s, μ_m, σ_m^2, σ_{sm}, σ_{s-m}^2 will all have the order of magnitude dt; this holds both in the usual theory of diffusion processes and in the game-theoretic framework (for a partial explanation, see Shafer and Vovk 2001, Chapter 9). So the right-hand side of (21) will not exceed $O\left((dt)^{3/2} + \frac{dt}{T}\right) + \frac{\epsilon\sigma_{s-m}^2}{2}$. For small enough ϵ, this should be much less than dt, the typical order of magnitude for μ_s, μ_m, σ_m^2, and σ_{sm}.

The data we consider in Section V are only monthly and cover only a few decades, and so they do not allow us to achieve the happy results suggested by these extreme asymptotics. In fact, the tightest bounds we can achieve with these data occur when we choose ϵ equal to 1.

B. The Theoretical Performance Deficit for Long Markets

The next proposition is a precise statement about the theoretical performance deficit $\sigma_{s-m}^2/2$.

Proposition 2 *For any* $\alpha \in (0,1]$ *and any* $\varepsilon \in (0,1]$, *the efficient market hypothesis for* m *predicts that*

$$\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m + \frac{1}{2}\sigma_{s-m}^2 < \frac{E_1}{\varepsilon} + E_2 + \frac{\ln\frac{1}{\alpha}}{N\varepsilon} + \frac{\varepsilon}{2}\sigma_{s-m}^2$$

at level α *in the long CAPG, where*

$$E_1 := \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(m_n) - \gamma((1-\varepsilon)m_n + \varepsilon s_n)\Big),$$

$$E_2 := \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(s_n) - \gamma(m_n)\Big),$$

and the functions Γ *and* γ *are defined in the statement of Proposition 1.*

This time we have broken the error stemming from the fundamental approximation into two parts. The first part, E_1/ε, usually increases as ε is made smaller, while the second part, E_2, is not affected by ε.

Again, we aim to choose α and ε so that α defines a reasonable efficient market hypothesis but the total error, in this case

$$\frac{E_1}{\varepsilon} + E_2 + \frac{\ln\frac{1}{\alpha}}{N\varepsilon} + \frac{\varepsilon}{2}\sigma_{s-m}^2, \tag{23}$$

is small. When this is achieved, the proposition says that $\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m + \frac{1}{2}\sigma_{s-m}^2 \lessapprox 0$, or

$$\frac{1}{N}\ln W_m - \frac{1}{N}\ln W_s \gtrapprox \frac{1}{2}\sigma_{s-m}^2.$$

In order for this validate the theoretical performance deficit $\sigma^2_{s-m}/2$ as a measure of s's performance, we need the error (23) to be small relative to all three terms in this approximate inequality. This evidently requires ε itself to be small. When $\varepsilon = 1$, (23) is larger than $\sigma^2_{s-m}/2$.

C. The Long-Short CAPM

Now we turn to the long-short case.

Proposition 3 *For any* $\varepsilon \in \left(0, \frac{\delta}{1+C}\right)$ *and* $\alpha \in (0,1]$, *the efficient market hypothesis for m predicts that*

$$\left|\mu_s - \mu_m + \sigma^2_m - \sigma_{sm}\right| < \frac{E}{\varepsilon} + \frac{\ln\frac{2}{\alpha}}{N\varepsilon} + \frac{\varepsilon}{2}\sigma^2_{s-m}$$

at level α *in the long-short CAPG with parameters C and* δ *(see p. 26), where*

$$E := \max_{j\in\{-1,1\}} \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(m_n) - \gamma((1-j\varepsilon)m_n + j\varepsilon s_n)\Big)$$

and the Γ *and* γ *are defined as in Proposition 1.*

D. The Theoretical Performance Deficit for Long-Short Markets

Proposition 4 *For any* $\alpha \in (0,1]$, *any parameters C and* δ *(see p. 26), and any* $\varepsilon \in \left(0, \frac{\delta}{1+C}\right)$, *the efficient market hypothesis for m predicts*

$$\left|\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m + \frac{1}{2}\sigma^2_{s-m}\right| < \frac{E_1}{\varepsilon} + E_2 + \frac{\ln\frac{2}{\alpha}}{N\varepsilon} + \frac{\varepsilon}{2}\sigma^2_{s-m}$$

at level α *in the long-short CAPG, where*

$$E_1 := \max_{j\in\{-1,1\}} \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(m_n) - \gamma((1-j\varepsilon)m_n + j\varepsilon s_n)\Big),$$

$$E_2 := \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(s_n) - \gamma(m_n)\Big) \bigvee \frac{1}{N}\sum_{n=1}^{N}\Big(\Gamma(m_n) - \gamma(s_n)\Big),$$

and the functions Γ *and* γ *are defined in the statement of Proposition 1.*

V. Some Empirical Examples

In this section, we check the game-theoretic CAPM's predictions against data on returns over the past three or four decades for a few well known stocks. We also investigate what the game-theoretic CAPM says about the equity premium by looking at two much longer sequences of returns for government and commercial bonds, one for the United States and one for Britain. All our tests use monthly data, with significance level $\alpha = 0.5$, corresponding to the hypothesis that SPECULATOR cannot do twice as well as the market index m, and mixing coefficient $\varepsilon = 1$.

Empirical tests of the classical CAPM do not emphasize returns on individual stocks. The classical CAPM cannot be tested at all until it is combined with additional hypotheses about the variability of individual securities, and in order to avoid putting the weight of a test on these additional hypotheses, one emphasizes portfolios, sometimes across entire industries, instead of individual securities. Moreover, even studies on returns from portfolios tend to be inconclusive, because of the substantial remaining variability orthogonal to the market and because the additional hypotheses still play a large role. Because the efficient market hypothesis is much weaker than the assumptions that go into the classical CAPM, we do not expect the game-theoretic CAPM to provide tighter bounds than the classical CAPM. So in order to find examples where the game-theoretic CAPM provides reasonably tight bounds on the relation between average return and volatility, or where the theoretical performance deficit provides an interesting bound on performance, we will probably need to look at large portfolios.

Moreover, the asymptotic analysis on p. 29 suggests that we will need to look at longer periods of time, perhaps with data sampled daily, in order to get tight bounds. And even a good understanding of how often the efficient market hypothesis at a given significance level is valid for individual securities in a given market would require a comprehensive and careful study, with due attention to survivorship bias and other biases.

This section should, however, make clear how our results can be applied to data. It shows the kinds of bounds that the game-theoretic CAPM can achieve with no assumptions beyond the level α for the efficient market hypothesis.

A. Twelve Stocks

The twelve stocks we now consider are listed in Table I. They are hardly a random or representative sample. We chose them because of their familiarity, and they are all still being traded. We did choose them, however, before making the calculations shown here; no other companies were chosen and then omitted because of the results they gave. Our data are from Yahoo, and they cover different time periods for different stocks, as indicated in the table. As the market index m, we use the S&P 500.

As we have already mentioned, we use $\alpha = 0.5$ and $\varepsilon = 1$. For all twelve companies, $\varepsilon = 1$ gives a better value for the bounds in the long-short case, both for Proposition 2 and Proposition 4, than any other $\varepsilon \in (0,1]$. (Strictly speaking, ε should satisfy $\varepsilon < \frac{\delta}{1+C}$ (see Proposition 3), but our results will be accurate if δ is close to 1 and C is close to 0.) So the inequalities we are checking are:

Long-Short CAPM Inequality (Proposition 3 on p. 31) with $\alpha = 0.5$ **and** $\varepsilon = 1$

$$\left|\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm}\right| < E + \frac{2\ln 2}{N} + \frac{1}{2}\sigma_{s-m}^2.$$

Long CAPM Inequality (Proposition 1 on p. 28) with $\alpha = 0.5$ **and** $\varepsilon = 1$

$$\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm} < E + \frac{\ln 2}{N} + \frac{1}{2}\sigma_{s-m}^2.$$

Table I
The Twelve Stocks

The number N, the number of monthly returns, is one less than the number of months for which we have data.

Company	Ticker Code	Time Period	N
IBM	IBM	January 1962–June 2001	473
General Electric	GE	January 1962–June 2001	473
Microsoft	MSFT	March 1986–June 2001	183
Boeing	BA	January 1970–June 2001	377
Du Pont (E.I.) de Nemours	DD	January 1970–June 2001	377
Consolidated Edison	ED	January 1970–June 2001	377
Eastman Kodak	EK	January 1970–June 2001	377
General Motors	GM	January 1970–June 2001	377
Procter and Gamble	PG	January 1970–June 2001	377
Sears/Roebuck	S	January 1970–June 2001	377
AT&T	T	January 1970–June 2001	377
Texaco	TX	January 1970–June 2001	377

(Remember that both E and N depend on the stock, and E is not the same in the long as in the long-short case.) In Table II, we report numerical values for these inequalities for our twelve companies.

Microsoft and Sears are the only stocks for which our bounds do not hold. Microsoft's spectacular performance violated both the long-short inequality (which reduces to $266.6 < 142.9$ in this case) and the long inequality (which reduces to $266.6 < 97.1$). Sears's impressive underperformance violated the long-short inequality ($68.0 < 65.5$) but not, of course, the long inequality ($-68.0 < 47.1$). Both performances can reasonably be put into the category of unusual events that could not have been anticipated, and for this reason the violations do not make us uncomfortable with our efficient market hypothesis. One other stock, AT&T, also falls outside our expectations. Our bounds hold for this stock, but the error in the fundamental approximation ($E = 217.3$) is so great that these bounds are uninteresting; this error seems to be the result of a fall in share price from 32.52 to 8.34 in a single month, February 1984.

Table II
CAPM Empirical Results

This table gives numerical values for the twelve stocks discussed in Subsection A and the two bond series discussed in Subsection B. In parentheses after the ticker code for the stock (or the abbreviation USA or UK) we give the ratio W_s/W_m, a direct measure of performance. All other numerical values are in basis points (1bp = 0.0001). The numbers *not* in parentheses are for the long-short inequality (Proposition 3). LHS, with the minus signs removed, is the left-hand side of that inequality; RHS is the right-hand side. For the long inequality (Proposition 1), use the values in parentheses for RHS and E, halve the value of $\frac{2\ln 2}{N}$, and restore the minus sign on the LHS.

Code (W_s/W_m)	μ_s	μ_m	σ_m^2	σ_{sm}	LHS	RHS	E	$\frac{1}{N}2\ln 2$	$\frac{1}{2}\sigma_{s-m}^2$
IBM (0.23)	83.1	99.4	19.6	18.7	−15.4	46.2 (31.4)	0.6 (0.4)	29.3	16.3
GE (0.91)	108.9	99.4	19.6	22.6	6.4	38.6 (23.5)	0.4 (0.0)	29.3	8.9
MSFT (50.3)	401.8	121.7	21.6	35.1	266.6	142.9 (97.1)	8.1 (0.2)	75.8	59.0
BA (1.7)	161.8	110.6	21.2	24.1	48.2	76.6 (54.9)	4.4 (1.1)	36.8	35.4
DD (0.76)	116.9	110.6	21.2	20.1	7.4	53.0 (33.9)	1.2 (0.4)	36.8	15.0
ED (1.3)	133.2	110.6	21.2	10.3	33.5	77.0 (54.9)	15.4 (11.7)	36.8	24.8
EK (0.12)	67.2	110.6	21.2	16.2	−38.4	56.7 (38.3)	1.6 (1.6)	36.8	18.4
GM (0.32)	97.6	110.6	21.2	19.0	−10.8	57.6 (39.2)	1.1 (1.1)	36.8	19.7
PG (0.43)	97.3	110.6	21.2	15.3	−7.4	53.1 (34.8)	1.5 (1.5)	36.8	14.8
S (0.032)	42.2	110.6	21.2	20.7	−68.0	65.5 (47.1)	6.4 (6.4)	36.8	22.3
T (0.016)	31.1	110.6	21.2	15.3	−73.6	280.8 (262.5)	217.3 (217.3)	36.8	26.7
TX (0.84)	117.2	110.6	21.2	14.9	12.9	56.6 (36.5)	1.5 (−0.1)	36.8	18.2
USA (0.0031)	36.2	84.7	23.7	0.3	−25.1	35.6 (16.4)	15.0 (0.3)	8.9	11.7
UK (0.0097)	41.7	55.7	11.3	4.5	−7.3	13.2 (8.0)	5.4 (0.2)	3.8	5.9

The information in Table II is sufficient to enable the reader to calculate our bounds for any other significance level α and for any other mixing coefficient ε.

Now we consider how well the twelve companies satisfy the bounds our theory gives for the theoretical performance deficit (TPD). In this case, we are considering the following inequalities:

Long-Short TPD Inequality (Proposition 4 on p. 31) with $\alpha = 0.5$ and $\varepsilon = 1$

$$\left|\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m + \frac{1}{2}\sigma_{s-m}^2\right| < E_1 + E_2 + \frac{2\ln 2}{N} + \frac{1}{2}\sigma_{s-m}^2.$$

Table III
TPD Empirical Results

The same conventions are used as in the preceding table.

Code	$\frac{1}{N}\ln W_s$	$\frac{1}{N}\ln W_m$	$\frac{1}{2}\sigma^2_{s-m}$	LHS	RHS	E_1+E_2	$\frac{1}{N}2\ln 2$	$\frac{1}{2}\sigma^2_{s-m}$
IBM	58.3	89.6	16.3	−15.0	47.0 (32.2)	1.4 (1.2)	29.3	16.3
GE	87.7	89.6	8.9	7.0	39.4 (24.4)	1.3 (0.8)	29.3	8.9
MSFT	324.9	110.8	59.0	273.2	152.4 (106.6)	17.6 (9.7)	75.8	59.0
BA	114.7	100.1	35.4	50.0	79.6 (57.9)	7.4 (4.1)	36.8	35.4
DD	92.8	100.1	15.0	7.7	53.8 (34.6)	2.0 (1.2)	36.8	15.0
ED	107.7	100.1	24.8	32.5	88.7 (55.3)	27.1 (12.1)	36.8	24.8
EK	43.0	100.1	18.4	−38.7	58.2 (38.6)	3.1 (1.8)	36.8	18.4
GM	69.6	100.1	19.7	−10.7	58.7 (39.9)	2.2 (1.8)	36.8	19.7
PG	77.6	100.1	14.8	−7.7	54.7 (34.9)	3.1 (1.7)	36.8	14.8
S	8.4	100.1	22.3	−69.3	71.8 (46.7)	12.8 (6.0)	36.8	22.3
T	−9.4	100.1	26.7	−82.7	498.2 (259.4)	434.7 (214.3)	36.8	26.7
TX	95.6	100.1	18.2	13.8	57.9 (37.9)	2.9 (1.2)	36.8	18.2
USA	36.1	73.0	11.7	−25.3	35.9 (16.8)	15.4 (0.7)	8.9	11.7
UK	37.4	50.2	6.0	−6.9	16.0 (8.9)	6.3 (1.1)	3.8	5.9

Long TPD Inequality (Proposition 2 on p. 30) with $\alpha = 0.5$ **and** $\varepsilon = 1$

$$\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m + \frac{1}{2}\sigma^2_{s-m} < E_1 + E_2 + \frac{\ln 2}{N} + \frac{1}{2}\sigma^2_{s-m}.$$

As we noted in the preceding section, we do not get interesting bounds for the theoretical performance deficit when $\varepsilon = 1$, and this is evident here from the fact that the deficit appears on both sides of each inequality. We can nevertheless look at the bounds as a check on our efficient market hypothesis. Table III reports results for the same companies and the same periods as Table II. On the whole, the numbers are quite close to those in Table II, but this time only Microsoft violates the bounds.

Table IV
The Two Bond Series

These two series of bond returns were obtained from Global Financial Data (GFD). In the case of the British series, the GFD files listed were supplemented by earlier data provided to us directly by GFD's Bryan Taylor. Once again, the number N of monthly returns is one less than the number of months for which we have data.

		American Series	British Series
Code		USA	UK
Time Period		January 1871–June 2001	June 1700–June 2001
N		1565	3612
s	Name	USA Total Return Commercial/T-bill Index	United Kingdom 10-year Government Bond Total Return Index
	GFD File	TRUSABIM.csv	TRGBRGVM.csv
m	Name	S&P 500 Composite Total Return Index	UK FT-Actuaries All-Share Total Return Index
	GFD File	TRSPXM.csv	_TFTASM.csv

B. The Equity Premium

We can apply the game-theoretic CAPM to bonds as well as to stocks, provided that we adopt the appropriate efficient market hypothesis: a speculator cannot substantially beat the index m when he is allowed to hold (and also short, if we want to apply the long-short CAPM) both m and the bonds. We now report on the results of applying the model to two series of bond returns, one American and one British. The sources of our data are listed in Table IV, and our results are summarized in the last two rows of Tables II and III.

The numbers in Table II confirm the usual result that bonds fall below the security market line: $\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm}$ is negative for both series. For the American series, this CAPM deficit is -25.1bp, and for the British series, it is -7.3bp. Because these values are negative, the long CAPM is automatically satisfied.

The long-short CAPM is also satisfied for both series ($25.1 < 35.6$ and $7.3 < 13.2$). These are only the results for $\alpha = 0.5$ and $\epsilon = 1$, but because the error E in the fundamental approximation is substantial, changing these parameters will not make much difference. Table III shows that the TPD inequalities are also satisfied by both series.

The amount by which bonds fall below the security market line can be thought of as the *equity premium*—a premium paid for holding equity that has the same covariance with the market as bond. Our results show that an equity premium may exist but that it remains within the bounds of our model and hence does not violate our efficient market hypothesis.

VI. Discussion

The purpose of this article is to introduce, as clearly as possible, a completely game-theoretic capital asset pricing model, derived from an efficient market hypothesis alone, with no assumptions about the beliefs or preferences of investors. The efficient market hypothesis is very weak; as we have seen, it really says only that a speculator cannot beat the market by a substantial factor using some rather obvious strategies. The predictions of this game-theoretic model are loose, in the sense that they give fairly wide bounds on the relation between average return and covariance with the market. But these bounds are themselves quite precise and can therefore be tested with no auxiliary stochastic model for individual returns. Because of this precision, the new game-theoretic model compares favorably with the classical CAPM, raising the question of whether the classical model's much stronger assumptions give it any greater predictive power. It also raises the question of whether the established generalizations of the classical CAPM, which also rely on stochastic assumptions and on assumptions about the beliefs and preferences of investors, are headed in the right direction. The game-theoretic approach, which begins with an appropriate recognition of the role of speculation in governing the market, may be more fruitful.

This is an almost purely theoretical article. It introduces a new way of understanding capital asset pricing, with a number of features that require detailed explanation, and this has left little room for empirical validation and elaboration. We have provided some examples to show

how the new approach applies to data, but we have not made a serious attempt to demonstrate its empirical usefulness. Nor have we explored the implications of the new approach for the many ways that the classical CAPM is presently used in finance.

Game theory is very flexible, even notoriously so, and the game we have presented in this article can obviously be varied in many ways. Some variations may lead to interesting alternative relations between return and volatility. As a simple example, we may cite the game in §15.4 of Shafer and Vovk (2001), which leads to a model similar to one in this article but more primitive.

One of the most important features of the game-theoretic CAPM is the way it casts doubt on the classical assumption that risk is measured by volatility. The tradeoff between return and covariance demonstrated by the game-theoretic CAPM does not depend on this assumption. This suggests that risk might be measured in other ways within a pricing game. This might give us a better understanding of the equity premium and might produce tighter and more useful models relating return, volatility, and risk.

References

Campbell, John Y., 2000, Asset pricing at the millennium, *Journal of Finance* 55, 1515–1567.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets*. (Princeton University Press, Princeton, NJ).

Cochrane, John H., 2001, *Asset Pricing*. (Princeton University Press, Princeton, NJ).

Copeland, Thomas E., and J. Fred Weston, 1988, *Financial Theory and Corporate Policy*. (Addison-Wesley, Reading, MA) third edn.

Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.

Shafer, Glenn, and Vladimir Vovk, 2001, *Probability and Finance: It's Only a Game!* (Wiley, New York).

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.

Appendix A. Lower and Upper Probability

For simplicity, our exposition of the game-theoretic CAPM in the body of this article avoids any concept of probability. The asset pricing games we have studied provide, however, a convenient framework for the game-theoretic concepts of lower and upper probability used in Shafer and Vovk (2001).

Recall the interpretation of the level α at which the efficient market hypothesis predicts an event A: the smaller α the more emphatic the prediction. If we want to measure our belief that A will happen on a scale from 0 to 1, then it is natural to code prediction at level α as belief of strength $1-\alpha$. For example, when A is predicted at the 5% level, we might claim a 95% belief in A's happening.

We may use the term "lower probability" for the strength of belief in this sense. Formally, we call the number $\underline{\mathbb{P}}A$ defined by

$$\underline{\mathbb{P}}A := 1-\alpha_A,$$

where

$$\alpha_A := \inf\{\alpha \mid 0<\alpha\le 1 \text{ and the efficient market hypothesis predicts } A \text{ at level } \alpha\}, \qquad \text{(A1)}$$

the *lower probability* for A. Roughly speaking (neglecting the fact that the infimum in (A1) might not be attained), the lower probability of A is the degree of belief corresponding to the smallest α such that A is predicted at level α. Because SPECULATOR has a winning strategy for the goal A when $\alpha = 1$ (buy and hold security 0), we always have $0 \le \alpha_A \le 1$ and hence $0 \le \underline{\mathbb{P}}A \le 1$.

Writing A^{c} for A's complement ($\Omega \setminus A$), we set

$$\overline{\mathbb{P}}A := 1-\underline{\mathbb{P}}A^{\mathrm{c}},$$

and we call $\overline{\mathbb{P}}A$ the *upper probability* of A. This quantity measures how plausible A is—the degree to which there is no particular reason for believing its complement A^{c}. It can be shown, using the fact that SPECULATOR cannot make money for sure in the game, that $\underline{\mathbb{P}}A \le \overline{\mathbb{P}}A$ (see Shafer and Vovk 2001, pp. 14–15). The concepts of lower and upper probability apply to all the games studied in this article—the basic CAPG, the long CAPG, and the long-short CAPG.

The fact that our upper and lower probabilities are measured on a probability-like scale (zero to one) does not mean that they should be interpreted as frequencies. On the contrary, they can only be interpreted in terms of the strength with which we assert the efficient-market hypothesis for the particular index m in the particular securities market. If we expect that a strategy selected in advance will not do 25% better than m, then we will regard $1-\frac{1}{1.25}$, or 0.2, as a high lower probability. If we would not be surprised to see a strategy selected in advance beat m by a factor of 10, then we will regard $1-\frac{1}{10}$, or 0.9, as a low lower probability.

We can restate Propositions 1–4 in terms of lower probability as follows:

Corollary 1 *In the long CAPG,*

$$\underline{\mathbb{P}}\left\{\mu_s-\mu_m+\sigma_m^2-\sigma_{sm}<\frac{E}{\varepsilon}+\frac{\ln\frac{1}{\alpha}}{N\varepsilon}+\frac{\varepsilon\sigma_{s-m}^2}{2}\right\}\geq 1-\alpha$$

and

$$\underline{\mathbb{P}}\left\{\frac{1}{N}\ln W_s-\frac{1}{N}\ln W_m+\frac{1}{2}\sigma_{s-m}^2<\frac{E_1}{\varepsilon}+E_2+\frac{\ln\frac{1}{\alpha}}{N\varepsilon}+\frac{\varepsilon}{2}\sigma_{s-m}^2\right\}\geq 1-\alpha,$$

for any $\alpha\in(0,1]$ and any $\varepsilon\in(0,1]$, where E, E_1, and E_2 are defined as in the statements of Propositions 1 and 2. In the long-short CAPG with parameters C and δ,

$$\underline{\mathbb{P}}\left\{\left|\mu_s-\mu_m+\sigma_m^2-\sigma_{sm}\right|<\frac{E}{\varepsilon}+\frac{\ln\frac{2}{\alpha}}{N\varepsilon}+\frac{\varepsilon\sigma_{s-m}^2}{2}\right\}\geq 1-\alpha$$

and

$$\underline{\mathbb{P}}\left\{\left|\frac{1}{N}\ln W_s-\frac{1}{N}\ln W_m+\frac{1}{2}\sigma_{s-m}^2\right|<\frac{E_1}{\varepsilon}+E_2+\frac{\ln\frac{2}{\alpha}}{N\varepsilon}+\frac{\varepsilon}{2}\sigma_{s-m}^2\right\}\geq 1-\alpha,$$

for any $\alpha\in(0,1]$ and any $\varepsilon\in\left(0,\frac{\delta}{1+C}\right)$, where E, E_1, and E_2 are defined as in the statements of Propositions 3 and 4.

The four displayed inequalities resemble probability statements. In fact, they are statements about probability in the sense of Shafer and Vovk. We need to remember, however, that they say exactly the same thing as Propositions 1–4: SPECULATOR has a winning strategy in certain games.

Appendix B. Proofs

Proof of Proposition 1: First we study the accuracy of the fundamental approximation. When $x > -1$, we can expand $\ln(1+x)$ in a Taylor's series with remainder:

$$\ln(1+x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3\frac{1}{(1+\theta x)^3}, \tag{B2}$$

where θ, which depends on x, satisfies $0 \le \theta \le 1$. Since

$$\gamma(x) \le \frac{1}{3}x^3\frac{1}{(1+\theta x)^3} \le \Gamma(x),$$

we can see that (B2) implies

$$\ln(1+x) \le x - \frac{1}{2}x^2 + \Gamma(x) \tag{B3}$$

and

$$\ln(1+x) \ge x - \frac{1}{2}x^2 + \gamma(x). \tag{B4}$$

Notice that the functions Γ and γ are monotonically increasing.

SPECULATOR has a trivial winning strategy in the long CAPG with any significance level α and the auxiliary goal

$$\prod_{n=1}^{N}(1+\varepsilon s_n + (1-\varepsilon)m_n) < \frac{1}{\alpha}\prod_{n=1}^{N}(1+m_n).$$

On each round, he invests ε of his capital in s and $1-\varepsilon$ of his capital in m. But we can rewrite this auxiliary goal as

$$\sum_{n=1}^{N}\Big(\ln(1+\varepsilon s_n + (1-\varepsilon)m_n) - \ln(1+m_n)\Big) < \ln\frac{1}{\alpha}.$$

This implies

$$\sum_{n=1}^{N}\Big(\varepsilon s_n + (1-\varepsilon)m_n - \frac{1}{2}\varepsilon^2 s_n^2 - \frac{1}{2}(1-\varepsilon)^2 m_n^2 - \varepsilon(1-\varepsilon)s_n m_n$$
$$+\gamma(\varepsilon s_n + (1-\varepsilon)m_n) - m_n + \frac{1}{2}m_n^2 - \Gamma(m_n)\Big) < \ln\frac{1}{\alpha}$$

or

$$\sum_{n=1}^{N}\Bigg(\varepsilon\left(s_n - m_n + m_n^2 - s_n m_n\right) - \frac{1}{2}\varepsilon^2\left(s_n^2 + m_n^2 - 2 s_n m_n\right)$$
$$+\gamma(\varepsilon s_n + (1-\varepsilon)m_n) - \Gamma(m_n)\Bigg) < \ln\frac{1}{\alpha}.$$

This completes the proof. ▮

Proof of Proposition 2: Add the equality

$$\frac{\sigma_{s-m}^2}{2} = \frac{\sigma_s^2}{2} + \frac{\sigma_m^2}{2} - \sigma_{sm} \tag{B5}$$

and the inequality

$$\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m$$
$$\leq \left(\mu_s - \frac{1}{2}\sigma_s^2 + \frac{1}{N}\sum_{n=1}^{N}\Gamma(s_n)\right) - \left(\mu_m - \frac{1}{2}\sigma_m^2 + \frac{1}{N}\sum_{n=1}^{N}\gamma(m_n)\right)$$

to the inequality in Proposition 1. ▮

Proof of Proposition 3: We know from Proposition 1 (see Corollary 1) that

$$\underline{\mathbb{P}}\left\{\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm} < \frac{E}{\varepsilon} + \frac{\ln\frac{2}{\alpha}}{N\varepsilon} + \frac{\varepsilon\sigma_{s-m}^2}{2}\right\} \geq 1 - \frac{\alpha}{2}; \tag{B6}$$

therefore, it remains to prove

$$\underline{\mathbb{P}}\left\{\mu_s - \mu_m + \sigma_m^2 - \sigma_{sm} > \frac{E}{-\varepsilon} + \frac{\ln\frac{2}{\alpha}}{-N\varepsilon} + \frac{-\varepsilon\sigma_{s-m}^2}{2}\right\} \geq 1 - \frac{\alpha}{2}. \tag{B7}$$

(Proposition 3 follows from (B6), (B7), and the inequality

$$\underline{\mathbb{P}}(A \cap B) \geq \underline{\mathbb{P}}(A) + \underline{\mathbb{P}}(B) - 1$$

from Shafer and Vovk 2001, Proposition 8.10.3 on p. 186.)

Consider a strategy for SPECULATOR in the long-short CAPG that calls for investing $-\varepsilon$ of his capital in s and investing $1+\varepsilon$ of his capital in m on every round. This strategy's return on round n is

$$-\varepsilon s_n + (1+\varepsilon)m_n \geq -\varepsilon C + (1+\varepsilon)(-1+\delta) = -1+\delta+\varepsilon(-C-1+\delta) > -1+\delta+\frac{\delta}{1+C}(-C-1) = -1$$

(remember that $\varepsilon < \frac{\delta}{1+C}$) and so it does not risk bankruptcy for SPECULATOR. It wins the game with the significance level $\frac{\alpha}{2}$ and auxiliary goal

$$\prod_{n=1}^{N}(1-\varepsilon s_n + (1+\varepsilon)m_n) < \frac{2}{\alpha}\prod_{n=1}^{N}(1+m_n),$$

and the auxiliary goal can be transformed as follows (this is similar to the calculations in the proof of Proposition 1, with ε replaced by $-\varepsilon$):

$$\sum_{n=1}^{N}\left(\ln(1-\varepsilon s_n + (1+\varepsilon)m_n) - \ln(1+m_n)\right) < \ln\frac{2}{\alpha};$$

$$\sum_{n=1}^{N}\Bigg(-\varepsilon s_n + (1+\varepsilon)m_n - \frac{1}{2}\varepsilon^2 s_n^2 - \frac{1}{2}(1+\varepsilon)^2 m_n^2 + \varepsilon(1+\varepsilon)s_n m_n$$
$$+\gamma(-\varepsilon s_n + (1+\varepsilon)m_n) - m_n + \frac{1}{2}m_n^2 - \Gamma(m_n)\Bigg) < \ln\frac{2}{\alpha};$$

$$\sum_{n=1}^{N}\Bigg(-\varepsilon\left(s_n - m_n + m_n^2 - s_n m_n\right) - \frac{1}{2}\varepsilon^2\left(s_n^2 + m_n^2 - 2s_n m_n\right)$$
$$+\gamma((1+\varepsilon)m_n - \varepsilon s_n) - \Gamma(m_n)\Bigg) < \ln\frac{2}{\alpha}.$$

This proves (B7). ∎

Proof of Proposition 4: One half of this proposition is Proposition 2 (with α replaced by $\frac{\alpha}{2}$), and the other half is obtained by adding (B5) and

$$\frac{1}{N}\ln W_s - \frac{1}{N}\ln W_m$$
$$\geq \left(\mu_s - \frac{1}{2}\sigma_s^2 + \frac{1}{N}\sum_{n=1}^{N}\gamma(s_n)\right) - \left(\mu_m - \frac{1}{2}\sigma_m^2 + \frac{1}{N}\sum_{n=1}^{N}\Gamma(m_n)\right)$$

to the inner inequality in (B7). ∎

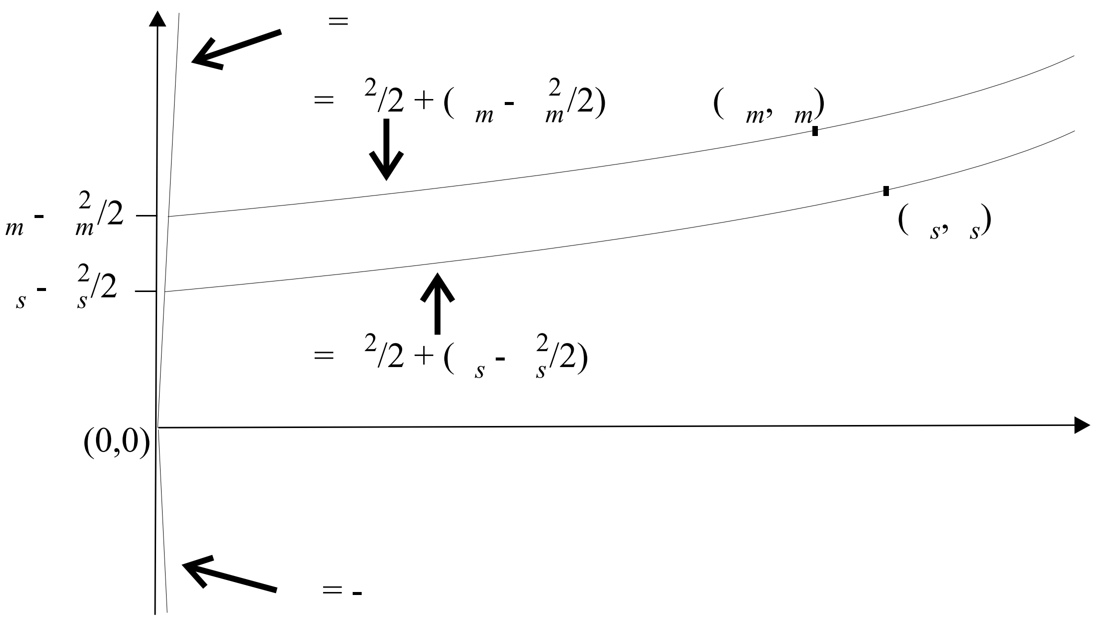

Figure 1. Indifference curves in the (σ,μ)**-plane.** Each curve is a parabola of the form $\mu = \frac{1}{2}\sigma^2 + c$ for some constant c. A speculator who is concerned only with final wealth will be approximately indifferent between two portfolios whose volatility-return pairs lie on the same such parabola. This figure also illustrates two additional points: (1) The indifference curve on which the market index m lies is called the *capital market parabola (CMP)*. (2) Because the minimum uncentered volatility σ compatible with a positive average return μ is μ, the line $\mu = \sigma$ represents the left-most boundary of the indifference curves in the positive quadrant. This line appears almost vertical because μ and σ are measured on very different scales; for a typical pair (σ,μ), σ^2 and μ are of the same order of magnitude, and so σ is much larger than μ.

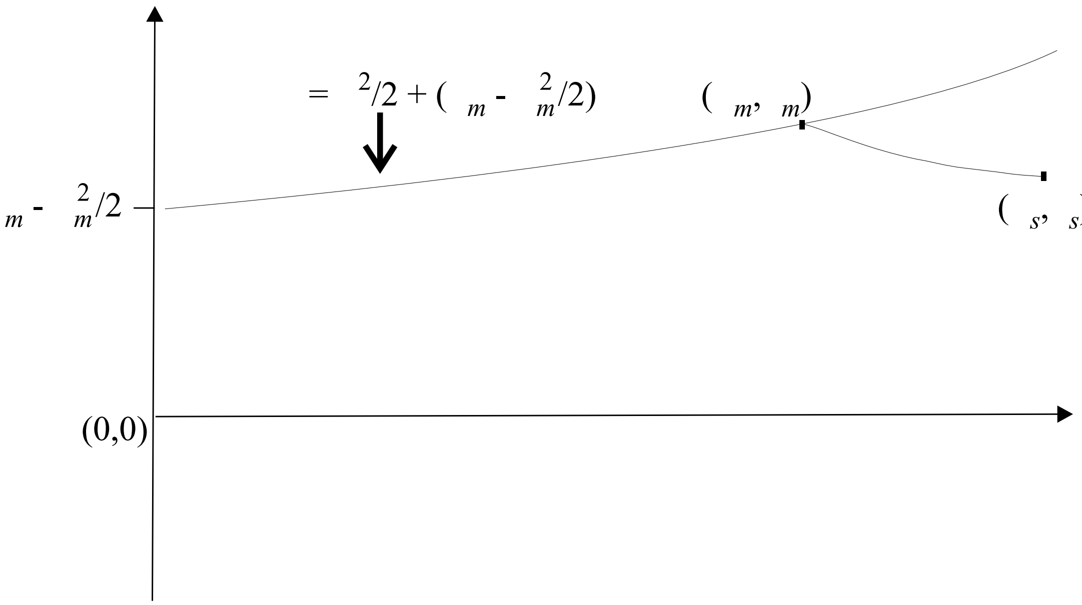

Figure 2. Mixing *m* with an underperforming portfolio *s*. The curve joining (σ_m, μ_m) and (σ_s, μ_s) is the trajectory traced by the volatility-return pair for the portfolio $\varepsilon s + (1-\varepsilon)m$ as ε varies from 0 to 1.

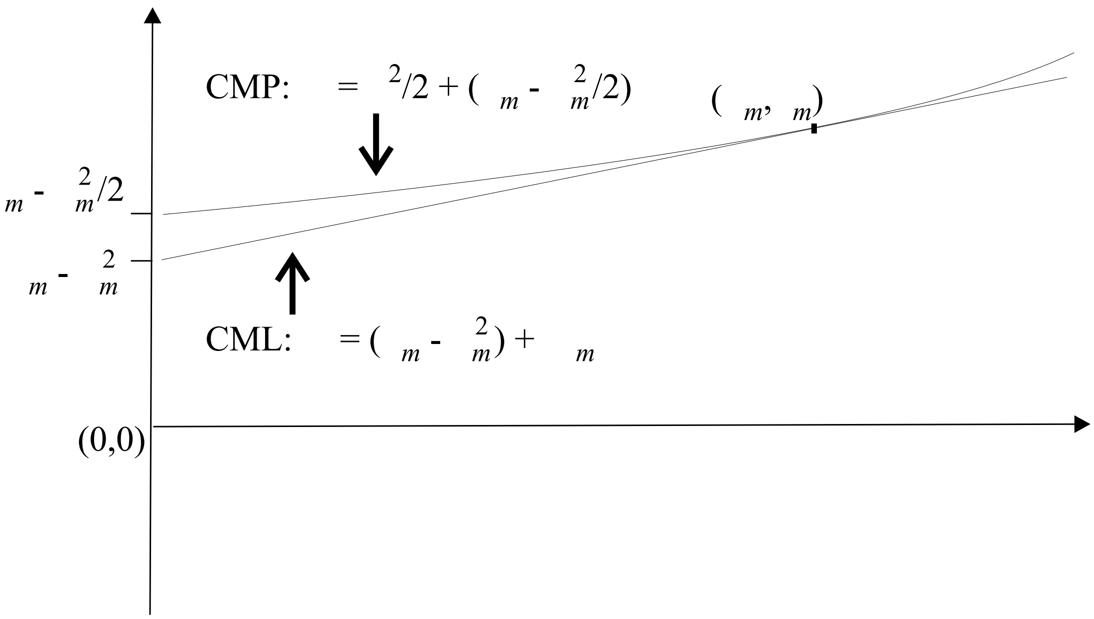

Figure 3. The capital market line (CML). This is the line tangent to the capital market parabola at (σ_m, μ_m). Our efficient market hypothesis implies that the volatility-return pair for a particular security or portfolio *s* should fall approximately on or below this line, even when a speculator cannot sell *s* short.

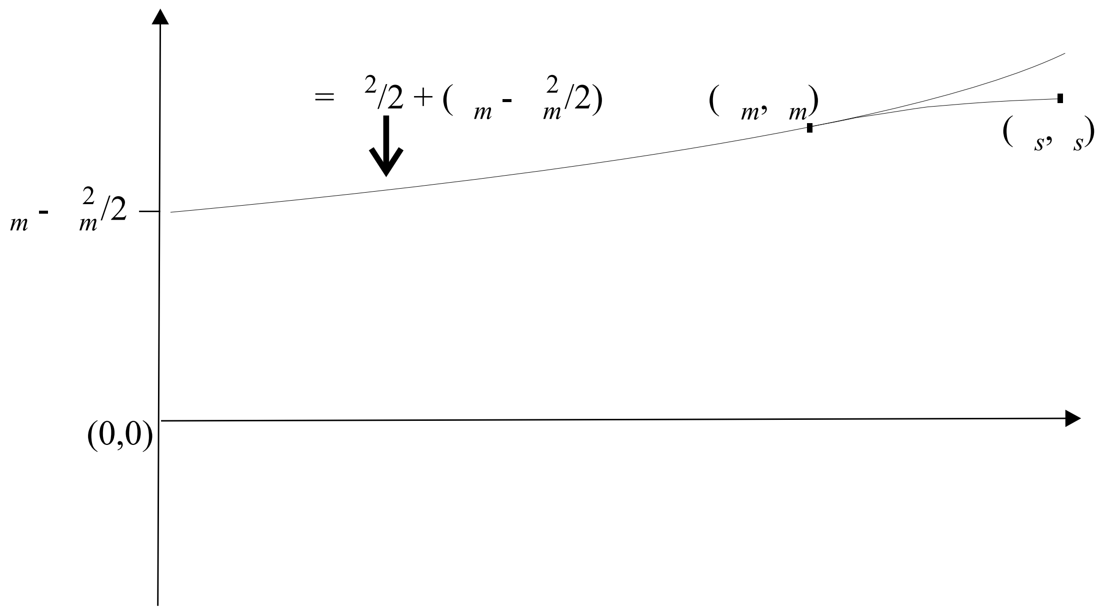

Figure 4. A trajectory for the long-short case. In the long-short case, the trajectory traced by the volatility-return pair for $\varepsilon s + (1-\varepsilon)m$ as ε varies from 0 to 1 must (1) approach (σ_m, μ_m) directly from the east, or (2) be tangent to the CMP (and therefore also to the CML) at (σ_m, μ_m). In this figure, it is tangent and approaches from the northeast. It could also approach from the southwest.