The risk and rewards of minimizing shortfall probability
Sid Browne
Journal of Portfolio Management; Summer 1999; 25, 4; ABI/INFORM Global
pg. 76

The Risk and Rewards of Minimizing Shortfall Probability

The risk may be worthwhile.

Sid Browne

Many different investment objectives and criteria have been suggested for choosing investment strategies. In a static setting, Markowitz [1952] suggests the mean-variance approach. Economic theory more formally postulates that an individual investor would choose an investment strategy to maximize expected utility of wealth and or consumption.

In other settings, other criteria might be more relevant. Rather than maximizing utility, investors in certain circumstances might be more concerned about *minimizing the probability of a shortfall,* where the shortfall is measured relative to a target return or a specific investment goal. Roy [1952] suggests this criterion in a static (one-period) framework and applies Chebyshev's inequality to obtain a criterion that is closely related to the mean-variance framework of Markowitz.

Investment strategies that minimize the probability of a shortfall can be more optimistically referred to as *probability-maximizing* strategies, in that they maximize the probability of reaching the investment goal. We consider the performance and implementation of such strategies in a dynamic multiperiod setting.

We show that dynamic probability-maximizing investment strategies have a variety of positive features that are attractive in a variety of economic settings, although there is also substantial investment risk. This latter point seems not that well understood by many practitioners, and an elaboration of this and other properties of probability-maximizing objectives is the main

SID BROWNE is vice president of firmwide risk at Goldman, Sachs and Co. in New York (NY 10005), and a professor of business at Columbia University in New York (NY 10027).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

objective of this article. (See Williams [1997] and Leibowitz, Bader, and Kogelman [1996].)

To illustrate our basic idea most directly, we adopt a simple continuous-time model, as in Black-Scholes [1973] or Merton [1971], where asset prices are lognormally distributed, and continuous trading is allowed. While our results hold for any number of assets, here we consider an economy with just two assets: a risky stock or equity index whose price at time t is denoted by S_t, and a riskless asset whose price at time t is denoted by B_t.

As in the Black-Scholes model, we assume that at a fixed time t, these prices can be written as[1]

$$S_t = S_0 \exp[(\mu - \sigma^2 / 2)t + \sigma\sqrt{t}Z] \qquad (1)$$

$$B_t = B_0 \exp(rt) \qquad (2)$$

where S_0 and B_0 are the initial prices, μ and σ are drift and volatility parameters of the risky stock, Z is a standard normal variable, and r is the instantaneous risk-free rate of interest. Note that the natural logarithm of the annual return relative of the stock has a normal distribution with mean $\mu - \sigma^2/2$ and variance σ^2.

For illustrative purposes, we first consider two different target returns: beating an *all-cash* strategy by 10%, and beating an *all-stock* strategy by 10%. (The choice of the level of 10%, while reasonable, is arbitrary and just for illustrative purposes. Other levels are considered later.) The all-cash strategy is of course the strategy that just buys and holds the riskless asset; similarly, the all-stock strategy just buys and holds the risky asset.

To illustrate some of the advantages and the possible pitfalls associated with shortfall probability minimization as an objective, we will compare its performance to the class of *constant allocations* (or mix) strategies, where the portfolio is continuously dynamically rebalanced so as to always hold a constant percent of wealth in the two assets. To enable exact calculations, we use the specific constant allocation strategy that maximizes a logarithmic utility function.[2]

In particular, if in the Black-Scholes model we take for illustrative purposes the numerical values $\mu = 0.15$, $\sigma = 0.3$, and $r = 0.07$, then the log-optimal constant allocation strategy always rebalances the portfolio so that 89% of wealth is held in the risky asset with the remaining 11% held in the risk-free asset, or cash.

As we will show, it takes ninety-eight years before the probability that this constant allocation strategy beats an all-cash strategy (by 10%) reaches 90%. Moreover, it will take 6,250 years for the probability that the constant mix strategy will outperform the stock itself (again by 10%), to reach 90%.

For the dynamic probability-maximizing strategy that we analyze, we will show that the 90% confidence or probability limits for beating the all-cash and the all-stock strategies by 10% fall to 0.04 years (fifteen days) and 2.60 years, respectively (see Browne [1997a] for a full mathematical treatment).

While these numbers for the probability-maximizing strategy are orders of magnitude better than the results of the constant mix strategy, there is an important dimension overlooked here: namely, that the probability-maximizing strategy is in general *riskier* than the constant allocation strategy. This can be illustrated most directly by noting that the *expected time* it takes for the constant allocation strategy to beat an all-cash strategy by 10% is 2.7 years, and the expected time for the constant mix strategy to beat an all-stock strategy is 172 years.

Although these numbers are not exactly a hearty endorsement of the constant mix strategy, they look quite good when compared with the results for the probability-maximizing strategy, since the expected time for a probability-maximizing strategy to beat an all-cash (or all-stock) strategy is *infinite*. This occurs because the dynamic probability-maximizing strategy is essentially *equivalent to the replicating or hedging strategy associated with a European digital option on the stock*, and thus entails the risk of not meeting the target return, or investment goal, by the associated time (see Browne [1997a]).[3]

This would occur if the stock finishes out of the money, and hence the digital option becomes worthless. We also show that there is another dimension of risk associated with probability maximization; there is a significant amount of *leveraging* necessary to undertake a probability-maximizing strategy relative to other strategies.

In some cases, probability maximization is entirely appropriate, and when used judiciously may even be risk-reducing. For the most part, this happens mainly when there is no finite deadline, or when the investor has an external source of income that can mitigate the investment risk associated with a probability-maximizing objective.[4]

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Even when probability maximization does entail more risk, our comparisons above show that this risk might be well worth taking, especially given that the implementation of a dynamic probability-maximizing strategy is completely equivalent to the static strategy that simply purchases a digital option. The implication of this equivalence is that an investor wishing to maximize the probability of achieving a preset target return or investment goal by a given time *need only purchase a (European) digital option* on the stock, with a specific exercise price determined by the underlying parameters, as we discuss below. This objective optimality property of options has important consequences for risk management.[5]

CONSTANT ALLOCATION PORTFOLIO STRATEGIES

In a constant allocation portfolio strategy, the portfolio holdings are continuously rebalanced so as to hold a constant percentage of wealth in the risky asset with the remaining percentage held in the risk-free asset. To use such a strategy, the investor needs to sell the risky asset as its price increases and buy it as its price falls. Such portfolio strategies are known to be optimal for a variety of economic investment objectives.[6]

For the model treated here, the wealth associated with a constant allocation portfolio strategy has a *lognormal* distribution, which allows exact probability calculations to be made. In particular, if θ is the rebalancing constant, then the portfolio holdings are continuously rebalanced so as to hold a constant θ of wealth in the risky asset and the remaining $1 - \theta$ in the risk-free asset.

The value of the portfolio, or wealth, associated with this constant allocation strategy at the fixed time t, $X_t(\theta)$, can be written as the lognormal random variable $X_t(\theta) = X_0 e^{N_1}$, where N_1 is a normal random variable with variance $\theta^2\sigma^2 t$ and mean $G(\theta)t$, where $G(\theta)$ is the quadratic function of the allocation fraction θ given by:[7]

$$G(\theta) = r + \theta(\mu - r) - \theta^2\sigma^2 / 2 \qquad (3)$$

The function $G(\theta)$ plays the role of the *continuous compounding rate* at which the wealth obtained from using strategy θ grows.

The choice of the rebalancing constant, θ, is determined in general by the investor's individual utility function and risk aversion parameter, and thus is investor-specific. To get a numerical value for θ, we use the value that maximizes the continuous compounding rate function $G(\theta)$. This value will be denoted by θ^*, and is given by

$$\theta^* = \frac{\mu - r}{\sigma^2} \qquad (4)$$

This value of θ maximizes the logarithmic utility function, which plays a central role in the continuous-time financial theory, and minimizes the expected time until a given target level of wealth is reached. The specific constant allocation strategy θ^* in (4) is often referred to as the *optimal growth* portfolio strategy. For the numerical values used here, $\mu = 15$, $\sigma = 0.3$, and $r = 0.07$, we obtain $\theta^* = 0.89$.[8]

Two other cases of alternative constant allocation strategies are of interest as comparative benchmarks: $\theta = 0$, and $\theta = 1$. In the first case, wealth is the portfolio associated with an all-cash strategy, and in the second case, wealth is the value of a portfolio fully invested in the stock itself.

The ratio of the wealth associated with any two constant allocation strategies is again lognormally distributed. The ratio of the wealth associated with the arbitrary constant allocation strategy θ to the wealth associated with the optimal growth strategy θ^*, $X_t(\theta)/X_t(\theta^*)$, can be written probabilistically, for any fixed t, as the lognormal random variable e^{-N_2}, where N_2 is a normal random variable with mean $\sigma^2(\theta^* - \theta)^2 t/2$ and variance $\sigma^2(\theta^* - \theta)^2 t$. This allows us to calculate some performance measures.[9]

Moreover, it can be shown that as t varies, the ratio process is a martingale, which has profound significance for the pricing of options and other contingent claims (this martingale property holds in much greater generality for the ratio of any portfolio strategy to the optimal growth strategy). In fact, the Black-Scholes formula for pricing contingent claims can be expressed solely in terms of the log-optimal wealth process.[10]

PROBABILITY CALCULATIONS

The fact that wealth under any constant mix strategy has a lognormal distribution enables some exact performance analysis via probability comparisons. In particular, for any given exceedence level ε, the probability that the growth-optimal strategy, θ^*, out-

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

performs any other constant allocation strategy θ by ε percent, by a *fixed* time T, is given explicitly by:[11]

$$\Phi\left(\frac{1}{2}M - \frac{\ln(1+\varepsilon)}{M}\right) \tag{5}$$

where Φ denotes the cumulative distribution function for a standard normal distribution, and where the value of M is given by

$$M = \sqrt{\sigma^2(\theta^* - \theta)^2 T} \tag{6}$$

The variable M is the standard deviation of the ratio of the different wealths at the time T. Observe that M is a monotonically increasing function of the time horizon T, and therefore increases without bound as T increases without bound. The probability given in Equation (5) thus tends to one as T increases without bound. That is, the probability that the growth-optimal strategy outperforms any other strategy θ by ε percent, for any given exceedence level ε, tends to one as the horizon increases without bound. This of course is one of the well known long-run optimality properties of the growth-optimal strategy. As observed by Rubinstein [1991], however, "the long run may be very long indeed."

Specifically, observe that for any given portfolio strategy θ, any given exceedence level ε, and any fixed given probability level $1 - \alpha$, the probability function in Equation (5) can be inverted in order to obtain the associated time T, which is given explicitly by:[12]

$$T = \left(\frac{q_\alpha + \sqrt{q_\alpha^2 + 2\ln(1+\varepsilon)}}{\sigma(\theta^* - \theta)}\right)^2 \tag{7}$$

where q_α is the $(1 - \alpha)$-th percentile or quantile of the standard normal distribution. Formally, q_α is defined as the root to the equation $\Phi(q_\alpha) = 1 - \alpha$, or equivalently, $q_\alpha = \Phi^{-1}(1 - \alpha)$, where Φ^{-1} is the inverse of the cumulative standard normal distribution function. Thus for example, the ninety-fifth percentile is denoted by $q_{0.05} = 1.645$.

Exhibit 1 shows the value of T, using an exceedence level of 10% (i.e., $\varepsilon = 0.1$), for various values of α and for the comparative strategies of all-cash and all-stock. The results are somewhat unsettling. It would take ninety-eight years for the optimal-growth strategy to have a 90% probability of beating an all-cash strategy, while it would take 6,250 years to be 90% certain of beating an all-stock strategy.

EXHIBIT 1
YEARS FOR OPTIMAL-GROWTH TO BEAT COMPETING STRATEGIES BY 10%

Probability $1 - \alpha$	Time Needed (in years) for θ^* to Beat θ by 10% All-Cash ($\theta = 0$)	All-Stock ($\theta = 1$)
0.900	98	6,250
0.950	158	10,080
0.990	310	19,824
0.999	542	34,720

The probability-maximizing strategy that we provide below will beat these numbers handily, but it incurs extra risk. For any exceedence level $\varepsilon > 0$, the *expected time* for the wealth of the optimal-growth strategy to beat the wealth associated with any other constant allocation strategy θ by ε percent is finite, and is given by:

$$\frac{2}{\sigma^2(\theta^* - \theta)^2}\ln(1+\varepsilon) \tag{8}$$

Evaluating (8) for $\varepsilon = 0.1$ with the illustrative numbers used previously gives the results in Exhibit 2.

DYNAMIC PROBABILITY-MAXIMIZING STRATEGIES

Suppose instead of following a constant allocation strategy, the investor dynamically changes the rebalancing level so as to minimize the probability of a shortfall by a given time. Equivalently, the investor wants to maximize the probability of reaching a given investment goal, say, b, by the given time T. The opti-

EXHIBIT 2
EXPECTED TIME FOR OPTIMAL-GROWTH TO BEAT COMPETING STRATEGIES BY 10%

Strategy	E (time to beat by 10%)
All-Cash ($\theta = 0$)	2.7 years
All-Stock ($\theta = 1$)	172.0 years

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

mal investment strategy for this problem in the simple Black-Scholes model considered here is equivalent to the replicating, or hedging, strategy of a European digital call option on the risky stock itself, with payoff b and some specific strike price. This equivalence follows from a direct economic argument based on pricing and valuation in a complete market.

Specifically, since the objective of the investor is only to reach the goal b, a probability maximizer would never choose a strategy for which terminal wealth would exceed b, since doing so would increase the cost of achieving that terminal wealth without increasing the associated rewards. Likewise, a probability-maximizing investor would never choose a strategy for which terminal wealth would give a value strictly between 0 and b, since the investor would gain resources by setting terminal wealth equal to 0 in that state.

Thus, the most efficient strategy for maximizing the probability of reaching the stated goal b is to purchase, or dynamically replicate, a digital option. The strike price of this optimal digital option is now obtained by equating the investor's initial wealth with the Black-Scholes value of a digital call option, and then solving for the strike price. (See Browne [1997a] for a complete analysis of related probability-maximizing strategies for more complicated models.)

Specifically, the value at time t of a digital option on a stock with price S_t, payoff b, and strike price K is, from the standard Black-Scholes equation, given by $C(t, S_t)$ where

$$C(t, S_t) = be^{-r(T-t)}\Phi(d_2) \tag{9}$$

where d_2 is given by

$$d_2 = \frac{\ln(S_t / K) + (r - \sigma^2 / 2)(T - t)}{\sigma\sqrt{T - t}} \tag{10}$$

Setting the initial price $C(0, S_0)$ equal to initial wealth X_0 gives the appropriate value for the strike price K for the probability-maximizing objective.[14]

The probability-maximizing strategy is thus completely equivalent to the dynamic replicating strategy for this particular digital option. This is the strategy that is continuously rebalanced so as to hold:[15]

$$\frac{1}{\sigma\sqrt{T - t}}\left[\frac{\phi(\nu_t)}{\Phi(\nu_t)}\right] \tag{11}$$

percent of wealth in the risky stock at time t with the remainder held in the risk-free asset, where ϕ and Φ denote the density and the cumulative distribution function of a standard normal variate, and where the quantity ν_t is the value of d_2 for this particular "probability-maximizing" strike price.

It turns out that ν_t is the quantile of the standard normal distribution associated with the *percentage of the goal reached by time t*. Explicitly, $\nu_t = \Phi^{-1}(z_t)$, where z_t, the percentage of the goal reached, is given by $z_t = x/[be^{-r(T-t)}]$, where x is the wealth level at time t.

Equation (11) represents the dynamic probability-maximizing strategy as the product of two separate positive effects. The first component is a purely time-dependent effect

$$\frac{1}{\sqrt{\sigma^2(T - t)}}$$

which increases (without bound) as the time until the deadline, T − t, decreases. The second quantity is a scalar determined solely by the percentage of the (effective) investment goal currently achieved, $\phi(\nu)/\Phi(\nu)$, which decreases as the ratio of current wealth to the effective goal, z, increases from 0 to 1.

Since it is equivalent to a digital option, the probability-maximizing strategy incurs the risk that with some probability the investment goal will not be reached, and the initial wealth X_0 will have been completely lost by time T. This possibility must be contrasted with the fact that under a constant allocation strategy, a complete loss of the initial stake is not possible in finite time in the continuous-time model. Nevertheless, the risk of a complete loss can be small enough over a particular horizon that the approach may be an attractive investment strategy. (We analyze the amount of *leverage* necessary to undertake such a strategy later.)

Maximizing the Probability of Beating Another Strategy

The probability-maximizing strategy can be generalized to treat a comparative rather than constant level b investment goal. We call this case *active probability maximization*, where performance is measured relative to another benchmark portfolio.

Suppose the investor wants to beat a given competing portfolio strategy by a given exceedence level ε (e.g., 10%). The competing portfolio strategy is

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

assumed to allocate θ_t% to the risky stock at time t, with θ_t known at each time t. For this problem, the corresponding active probability-maximizing strategy is no longer equivalent to a simple digital option, although the generalization is not hard to work out.[16]

If the competing strategy is a constant allocation strategy, with allocation constant θ, then it turns out that the best you can do against the given strategy θ (with respect to maximizing the probability of beating it by ε%) is given by

$$\Phi\left(\Phi^{-1}\left(\frac{1}{1+\varepsilon}\right)+\sqrt{\sigma^2(\theta^*-\theta)^2T}\right) \quad (12)$$

where Φ^{-1} denotes the appropriate quantile of the standard normal distribution.

For a given strategy θ and a given exceedence level ε, we may now set the maximal probability of (12) equal to a given probability level, say, $1-\alpha$, and then solve for the corresponding time, T, that is needed to achieve the given probability. The result is given explicitly by:[17]

$$T=\left[\frac{\Phi^{-1}(1-\alpha)-\Phi^{-1}(1/[1+\varepsilon])}{\sigma(\theta^*-\theta)}\right]^2 \quad (13)$$

where again the appropriate quantiles of the standard normal distribution are given by Φ^{-1}.

We may now evaluate this for various values of θ, ε, and α. By taking $\theta = 0$ we get the relevant time to beat an all-cash strategy, and by taking $\theta = 1$ we get the relevant time to beat an all-stock strategy.

Exhibit 3 presents results for an exceedence level of 10% with the numerical values $\mu = 0.15$, $r = 0.07$, and $\sigma = 0.3$. For $\varepsilon = 0.1$, $1/(1+\varepsilon) = 1/1.1 = 0.91$, and $\Phi^{-1}(0.91) = 1$.

As Exhibit 3 shows, the active probability-maximizing strategy gives results that are orders of magnitude better than the comparative results for the constant allocation (optimal-growth) strategy analyzed. The downside of course, is that under this strategy, the terminal value of the portfolio at time T has positive probability of being 0, as it is essentially an options strategy. As Exhibit 3 shows, however, these probabilities can be made arbitrarily small by increasing the time horizon.

EXHIBIT 3
YEARS FOR PROBABILITY-MAXIMIZING STRATEGY TO BEAT COMPETING STRATEGIES BY 10%

Probability $1-\alpha$	Time Needed (in years) to Beat θ by 10% All-Cash ($\theta = 0$)	All-Stock ($\theta = 1$)
0.900	0.04 (15 days)	2.6
0.950	1.35	86
0.990	14.00	884
0.999	43.00	2,772

Risk-Taking and Leveraging

There is a side to the story that Exhibit 3 does not show: the amount of leverage undertaken by a probability maximizer. Leveraging or borrowing takes place when the percentage of wealth invested in the risky asset exceeds unity. The constant allocation strategy has not entailed leveraging, since θ^* for our example is less than one. (Of course, if $\theta^* > 1$, then, unless otherwise constrained, the investor must always leverage.) For the probability-maximizing strategy, we see from Equation (11) that leveraging occurs when

$$\phi(\nu) / \Phi(\nu) \geq \sqrt{\tau}$$

where $\tau = \sigma^2(T-t)$, or *the risk-adjusted remaining time.*

The leveraging necessitated by a probability-maximizing strategy can be analyzed by looking at the *borrowing region,* denoted in Exhibit 4 by the curve $z^*(\tau)$. Borrowing is necessitated with τ risk-adjusted time remaining only if $z(\tau) < z^*(\tau)$, where $z(\tau)$ is the actual proportion of the goal attained with τ risk-adjusted time units to go, and $z^*(\tau)$ is the calculated root to the equation $\phi(\nu)/\Phi(\nu) = \sqrt{\tau}$.[18]

Some select values of of τ and $z^*(\tau)$ are given in Exhibit 5. As Exhibit 5 shows, if $\tau = 0.05$ risk-adjusted time units remain until the deadline, the probability-maximizing investor must borrow (at the risk-free rate r) unless the investor is already 88% of the way to the goal. As the remaining risk-adjusted time to go *increases,* a probability-maximizing investor needs to borrow only at lower percentages. For example, if there is $\tau = 1$ unit of time left to go, the investor will need to borrow unless wealth at that time is more than 38% of the way to the investment goal.

It is important to note that *increasing the risk fac-*

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

EXHIBIT 4
$z^*(\tau)$ PLOTTED AGAINST τ

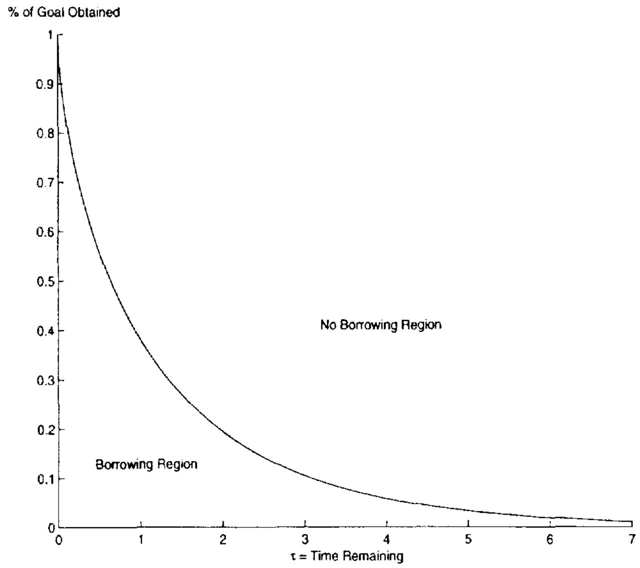

EXHIBIT 5
BORROWING REGION

τ	$z^*(\tau)$	τ	$z^*(\tau)$	τ	$z^*(\tau)$
0.001	0.99	0.50	0.56	1.00	0.38
0.050	0.88	0.55	0.54	1.50	0.27
0.100	0.82	0.60	0.51	2.00	0.19
0.150	0.77	0.65	0.49	2.50	0.14
0.200	0.73	0.70	0.48	3.00	0.10
0.250	0.70	0.75	0.46	3.50	0.08
0.300	0.67	0.80	0.44	4.00	0.06
0.350	0.64	0.85	0.43	4.50	0.04
0.400	0.61	0.90	0.41	5.00	0.03
0.450	0.58	0.95	0.39		

τ = risk-adjusted time to go; $z^*(\tau)$ = critical percentage of the distance to the goal attained with τ risk-adjusted time units remaining. Borrowing occurs if the actual $z(\tau)$ is less than $z^*(\tau)$.

tor of the stock, σ^2, has the same effect as increasing the actual time left to play, $T - t$. Therefore, for a higher risk factor, one would borrow less in the hopes of reaching the investment goal later.

CONCLUSIONS

We have provided a comprehensive analysis of the dynamic probability-maximizing strategy, and compared its performance to that of a constant allocation strategy. While probability maximization necessitates some extra risk, this risk might be worthwhile when we look at some of the performance disadvantages of the constant allocation strategy. We have also shown that it is quite simple to undertake a probability-maximizing strategy since it is completely equivalent to the purchase of a digital option, with a specific strike price.

In many scenarios, such as in corporate risk management settings, there is no finite deadline. A manager may rather be interested in minimizing the probability of *ever* going below a given shortfall level.

The major problem with probability maximization is that the payoff function is binary valued (1 at the investment goal and 0 elsewhere). Therefore, if there is a finite deadline, significant risk-taking occurs near the deadline if wealth is far from the investment goal. This is the case when there is no other source of income available to the investor.

If indeed there is an external source of income, this income stream can be used to mitigate the investment risk associated with a probability-maximizing objective. This contrasts with utility maximization objectives, where income is used to take *extra risk* by borrowing against future earnings.

ENDNOTES

[1]In Black and Scholes [1973] and Merton [1971], the price of the risky asset, S_t, is assumed to evolve according to the stochastic differential equation

$$dS_t = \mu S_t dt + \sigma S_t dW_t$$

where W_t is a standard Brownian motion, or Wiener process, and the price of the riskless asset, B_t, is assumed to evolve according to the differential equation: $dB_t = rB_t dt$. (A standard Brownian motion is a process with continuous sample paths that has independent increments. These increments are normally distributed with mean zero, and, in particular, for any t, W_t has a normal distribution with mean 0 and variance t.) The solution of the stochastic differential equation is given by the geometric Brownian motion

$$S_t = S_0 \exp[(\mu - \sigma^2/2)t + \sigma W_t]$$

Since W_t is normally distributed with mean 0 and variance t, for any fixed t this is probabilistically equivalent to the lognormal random variable given in Equation (1). Equation (2) is the solution to the differential equation solving for dB_t.

[2]See Perold and Sharpe [1988] for performance analysis of this and other strategies.

[3]Specifically, if S_t denotes the stock price in the Black-Scholes economy, a digital call option with *strike price K* and *payoff b* is an option that pays \$b to the bearer at time T if and only if the stock price exceeds the strike price at T, i.e., if $S_T \geq K$. Standard results on option pricing show that the (fair) price at time t for this option is

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$C(t, S_t) = be^{-r(T-t)}\Phi(d_2)$$

where d_2 is given by

$$d_2 = \frac{\ln(S_t / K) + (r - \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$

The associated hedging, or replicating, strategy is the dynamic trading strategy that holds $\partial C/\partial S_t$ shares of the stock at time t, with the remaining wealth held in the riskless asset.

[4]See, for example, Browne [1995, 1997b] for infinite horizons and Browne [1997a] for the finite-horizon problem with external income.

[5]In more general models with multiple stocks, the probability-maximizing strategy is still equivalent to a digital option, but on a particular index of the stocks, as discussed in Browne [1997a].

[6]In particular, with no further contributions or withdrawal of funds, a constant allocation strategy maximizes expected utility of terminal wealth for a concave utility function with constant relative risk aversion, as discussed in Merton [1971, 1990] or Hakansson [1970].

[7]In the Black-Scholes model, if θ is the constant fraction of wealth invested in the risky stock at time t, with the remainder held in the riskless asset, the total amount of money invested in the stock is $\theta X_t(\theta)$, and therefore the total number of shares held in the stock is $\theta X_t(\theta)/S_t$. Similarly, the total number of shares of the riskless asset held at time t is given by $(1 - \theta)X_t(\theta)/B_t$. The instantaneous change in wealth at time t can be written as

$$dX_t(\theta) = \theta X_t(\theta)\frac{dS_t}{S_t} + (1-\theta)X_t(\theta)\frac{dB_t}{B_t}$$

When dS_t and dB_t from endnote 1 are substituted into the equation, we obtain the stochastic differential equation

$$dX_t(\theta) = X_t(\theta)[r + \theta(\mu - r)]dt + X_t(\theta)\sigma\theta dW_t$$

whose solution is the geometric Brownian motion given by

$$X_t(\theta) = X_0\exp[G(\theta)t + \theta\sigma W_t]$$

where X_0 is the initial wealth level, and W_t is a standard Brownian motion. Since W_t is normally distributed with mean 0 and variance t, for any fixed t we can write this as

$$X_t(\theta) = X_0\exp[G(\theta)t + \theta\sigma\sqrt{t}Z]$$

where Z is a standard normal variable with mean 0 and variance 1.

[8]See Hakansson [1970], Merton [1990, Chapter 6], and Browne [1998] for further properties and analysis of the optimal growth and other constant allocation rules. Merton discusses the importance of the logarithmic utility function.

[9]Specifically, for any fixed t, we have

$$X_t(\theta) / X_t(\theta^*) = \exp\left[-\frac{\sigma^2}{2}(\theta^* - \theta)^2 t - \sigma(\theta^* - \theta)\sqrt{t}Z\right]$$

where Z is a standard normal variable. This follows directly from the last equation in endnote 7, since for any fixed t, we can write the ratio as

$$X_t(\theta) / X_t(\theta^*) = \exp\left[-[G(\theta^*) - G(\theta)]t - \sigma(\theta^* - \theta)\sqrt{t}Z\right]$$

where G is the function defined in Equation (3). Using Equation (3) and the definition of θ^* given in (4), we can write

$$G(\theta^*) = r + (\theta^*)^2\sigma^2/2$$

and, for an arbitrary θ, we can also write

$$G(\theta) = r + \theta\theta^*\sigma^2 - \theta^2\sigma^2/2$$

Thus:

$$G(\theta^*) - G(\theta) = (\theta^*)^2\sigma^2/2 + \theta^2\sigma^2/2 - \theta\theta^*\sigma^2$$

which is equivalent to $\sigma^2(\theta^* - \theta)^2/2$.

[10]Consider a derivative security that, for a given prespecified payoff function H(x), and a given prespecified deadline T, will pay the owner $H(S_T)$ at time T, where S_t is the price of the stock at time t.

Some standard examples of payoff functions are: 1) $H(x) = \max[0, x - K]$; 2) $H(x) = \max[0, K - x]$; and 3) $H(x) = b$ if $x \geq K$, 0 otherwise. These are the terminal payoff functions associated with 1) a European call option with strike price K; 2) a European put option with strike price K; and 3) a European digital option with payoff b and strike price K.

The "rational" or fair price at time t for this contingent claim can now be expressed, in terms of the log-optimal wealth $X_t(\theta^*)$, the stock $X_t(1)$, and the risk-free asset $X_t(0)$, as

$$X_t(0)E_t\left(\frac{H[X_T(1)]}{X_T(\theta^*)}\right)$$

where E_t denotes the conditional expectation on the relevant history at time t.

The Black-Scholes pricing formula of this equation holds in more general models of security prices and for much more general (sometimes path-dependent) payoff structures. For a more complete treatment, see Merton [1990].

[11]The probability that portfolio strategy θ^* beats portfolio strategy θ is

$$P[X_T(\theta^*) > (1 + \varepsilon)X_T(\theta)]$$

Some direct manipulation using the first equation in endnote 9 shows that this probability can be written as

$$P[\sigma(\theta^* - \theta)\sqrt{T}Z > -\frac{1}{2}\sigma^2(\theta^* - \theta)^2 T + \ln(1+\varepsilon)]$$

where Z has a standard normal distribution with mean 0 and variance 1, so $\sigma(\theta^* - \theta)\sqrt{T}Z$ has a normal distribution with mean 0 and variance $\sigma^2(\theta^* - \theta)^2T$. This probability is equal to the probability in Equation (5).

[12]Specifically, setting the probability in Equation (5) equal to $1 - \alpha$ gives

$$\Phi\left(\frac{1}{2}M - \frac{\ln(1+\varepsilon)}{M}\right) = 1 - \alpha$$

which can be inverted to

$$\frac{1}{2}M - \frac{\ln(1+\varepsilon)}{M} = \Phi^{-1}(1-\alpha)$$

Multiplying by M, subtracting appropriately, and using the standard normal quantile notation $q_\alpha = \Phi^{-1}(1 - \alpha)$, the value of M that achieves the probability of $1 - \alpha$ is given as the positive root to the quadratic equation

$$\frac{1}{2}M^2 - q_\alpha M - \ln(1+\varepsilon) = 0$$

This positive root is given by

$$M = q_\alpha + \sqrt{q_\alpha^2 + 2\ln(1+\varepsilon)}$$

When we substitute for M as it is defined in Equation (6), we get

$$\sqrt{\sigma^2(\theta^* - \theta)^2 T} = q_\alpha + \sqrt{q_\alpha^2 + 2\ln(1+\varepsilon)}$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

and finally when we square both sides and solve for T, we obtain Equation (7).

[13]The expected time formula in Equation (8) follows from the well-known fact that if Y_t is the geometric Brownian motion given by

$$Y_t = \exp[\gamma t + \beta W_t]$$

where W_t is a standard Brownian motion, then for any $\gamma > 0$, and any β, the expected time for Y to "grow $100 \times \lambda\%$" (i.e., the time until $Y_t = 1 + \lambda$) is given by

$$\frac{1}{\gamma}\ln(1+\lambda)$$

From the last equation in endnote 7 we observe that the relevant growth rate γ for a constant allocation strategy is the quadratic function $G(\theta)$ of Equation (3). It also follows that for any $\varepsilon > 0$, the expected time for the wealth of the optimal-growth strategy to beat the wealth associated with any other constant allocation strategy θ by ε is in fact given by Equation (8).

[14]The relevant strike price (K^*) is given by

$$K^* = S_0 \exp\left[(r - \sigma^2/2)T - \sigma\sqrt{T}\Phi^{-1}(X_0 e^{rT}/b)\right]$$

where S_0 denotes the initial stock price, and X_0 denotes the initial wealth level.

For more general models, the probability-maximizing strategy is again equivalent to the "replicating strategy" of a European digital call option, but now the option is on the *log-optimal wealth* for that model, rather than the stock itself. See Browne [1997a] for the details.

[15]The dynamic strategy in (11) is explained as follows. Observe that for a digital option, we have

$$S_t \frac{\partial}{\partial S_t} C(t,S_t)/C(t,S_t) = \frac{1}{\sigma\sqrt{T-t}}\left[\frac{\phi(d_2)}{\Phi(d_2)}\right]$$

Note further that d_2 of (10) evaluated at the strike price K^* as in endnote 14 is given by ν_t.

[16]As Browne [1997a] shows, the dynamic active probability-maximizing strategy invests

$$\theta_t + \frac{1}{\sqrt{\sigma^2(T-t)}}\left[\frac{\phi(\nu')}{\Phi(\nu')}\right]$$

% of its wealth in the risky stock at time t with the remainder invested in the riskless asset. The active probability-maximizing portfolio strategy decomposes into two additive parts. The first part mimics the tracking portfolio strategy θ_t, and the second part then overinvests in the risky stock according to the regular probability-maximizing strategy. The quantity ν' is given by $\Phi^{-1}(z')$ where z' is defined as the ratio of the probability-maximizing wealth to the competing portfolio strategy's wealth, divided by $1 + \varepsilon$; i.e., at time t the quantity z' is given by $w_t/(1 + \varepsilon)$, where w_t is the ratio of the probability-maximizing strategy to the wealth of the competing allocation strategy θ_t.

The maximal probability associated with the active probability-maximizing strategy for a competing constant allocation strategy θ is given by

$$\Phi\left(\Phi^{-1}\left(\frac{w}{1+\varepsilon}\right) + \sqrt{\sigma^2(\theta^* - \theta)^2(T-t)}\right)$$

where θ^* is the optimal-growth strategy of (3).

[17]Specifically, setting Equation (12) equal to $1 - \alpha$ gives

$$\Phi\left(\Phi^{-1}\left(\frac{1}{1+\varepsilon}\right) + \sqrt{\sigma^2(\theta^* - \theta)^2 T}\right) = 1 - \alpha$$

which can be inverted to

$$\Phi^{-1}\left(\frac{1}{1+\varepsilon}\right) + \sqrt{\sigma^2(\theta^* - \theta)^2 T} = \Phi^{-1}(1-\alpha)$$

Solving this last equation for T gives the result in Equation (13).

[18]If $\nu^*(\tau)$ denotes the unique root to the Equation $\phi(\nu)/\Phi(\nu) = \sqrt{\tau}$, then $z^*(\tau) = \Phi[\nu^*(\tau)]$. These roots are decreasing in τ, so borrowing occurs with τ risk-adjusted time remaining only when $z(\tau) < z^*(\tau)$, where $z(\tau)$ is the proportion of the goal attained with τ risk-adjusted time units to go.

REFERENCES

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637–659.

Browne, S. "Optimal Investment Policies for a Firm with a Random Risk Process: Exponential Utility and Minimizing the Probability of Ruin." *Mathematics of Operations Research*, 20 (1995), pp. 937–958.

——. "Reaching Goals by a Deadline: Digital Options and Continuous-Time Active Portfolio Management."*Proceedings of the 1997 IAFE Annual Conference*, 1997a. (To appear in *Advances in Applied Probability*, 1999.)

——. "Survival and Growth with a Fixed Liability: Optimal Portfolios in Continuous Time." *Mathematics of Operations Research*, 22 (1997b), pp. 468–493.

——. "The Return on Investment from Proportional Portfolio Strategies." *Advances in Applied Probability*, 30 (1998), pp. 216-238.

Hakansson, N.H. "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions." *Econometrica*, 38 (1970), pp. 587–607.

Leibowitz, M.L., L.N. Bader, and S. Kogelman. *Return Targets and Shortfall Risks*. Burr Ridge, IL: Irwin, 1996.

Markowitz, H.M. "Portfolio Selection."*Journal of Finance*, 7 (1952), pp. 77–91.

Merton, R. *Continuous Time Finance*. Cambridge: Blackwell, 1990.

——. "Optimum Consumption and Portfolio Rules in a Continuous Time Model." *Journal of Economic Theory*, 3 (1971), pp. 373–413.

Perold, A.F., and W.F. Sharpe. "Dynamic Strategies for Asset Allocation." *Financial Analysts Journal*, January-February 1988, pp. 16–27.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Roy, A. "Safety First and the Holding of Assets." *Econometrica*, 20 (1952), pp. 431–439.

Rubinstein, M. "Continuously Rebalanced Investment Strategies." *Journal of Portfolio Management*, Fall 1991, pp. 78–81.

Williams, J.O. "Maximizing the Probability of Achieving Investment Goals." *Journal of Portfolio Management*, Fall 1997, pp. 77–81.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.