JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS VOL. 39, NO. 1, MARCH 2004
COPYRIGHT 2004, SCHOOL OF BUSINESS ADMINISTRATION, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

Sharpe Ratios and Alphas in Continuous Time

Lars Tyge Nielsen and Maria Vassalou*

Abstract

This paper proposes modified versions of the Sharpe ratio and Jensen's alpha, which are appropriate in a simple continuous-time model. Both are derived from optimal portfolio selection. The modified Sharpe ratio equals the ordinary Sharpe ratio plus half of the volatility of the fund. The modified alpha also differs from the ordinary alpha by a second-moment adjustment. The modified and the ordinary Sharpe ratios may rank funds differently. In particular, if two funds have the same ordinary Sharpe ratio, then the one with the higher volatility will rank higher according to the modified Sharpe ratio. This is justified by the underlying dynamic portfolio theory. Unlike their discrete-time versions, the continuous-time performance measures take into account that it is optimal for investors to change the fractions of their wealth held in the fund vs. the riskless asset over time.

I. Introduction

This paper proposes and analyzes modified versions of the Sharpe ratio and Jensen's alpha, which are derived from optimal portfolio selection in a simple continuous-time model. The ordinary Sharpe ratio was proposed by Sharpe (1966) and is much used by practitioners. It is the ratio between expected or average excess return and risk, where risk is measured as standard deviation of return.

According to static mean-variance portfolio theory, if investors face an exclusive choice among a number of funds, then they can unambiguously rank them on the basis of their Sharpe ratios. A fund with a higher Sharpe ratio will enable all investors to achieve a higher expected utility. The *modified* or *instantaneous* Sharpe ratio is effectively the same as the discrete Sharpe ratio, except that the rates of return over finite time intervals are replaced by instantaneous rates of return.

We show that if investors face an exclusive choice among a number of funds, each of which has a constant instantaneous Sharpe ratio and if they are able to dynamically reallocate wealth between their chosen fund and a money market

*Nielsen, lars.nielsen@morganstanley.com, Morgan Stanley, 750 7th Avenue, New York, NY 10019, and Vassalou, maria.vassalou@columbia.edu, Graduate School of Business, Columbia University, 3022 Broadway, New York, NY 10027. The authors thank Zhiwu Chen (associate editor and referee) and especially Stephen J. Brown (the editor) for many useful comments. We are responsible for any errors. The views, thoughts and opinions expressed in this paper are those of the authors and should not in any way be attributed to Morgan Stanley or to Lars Tyge Nielsen as a representative, officer, or employee of Morgan Stanley.

account, then they can unambiguously rank the funds on the basis of their instantaneous Sharpe ratios. A fund with a higher instantaneous Sharpe ratio will enable all investors to achieve a higher expected utility.

The assumption of constant instantaneous Sharpe ratios is obviously restrictive, but it does allow the volatilities and expected excess returns of the funds to change stochastically over time. As long as a fund invests in an underlying portfolio that has a constant instantaneous Sharpe ratio, it may well engage in a dynamic strategy with respect to the fraction of asset value invested in the portfolio and the fraction invested in the riskless asset, or the degree of leverage employed. If the underlying portfolio has constant volatility, then the fund may also engage in a strategy that involves buying and selling contingent claims such as put and call options on the portfolio.

Even though our linkage of expected utility maximization and the instantaneous Sharpe ratio allows for dynamic strategies, it does not contradict the concerns about gaming of the Sharpe ratio expressed, for example, in Goetzman et al. (2002). That paper applies the Sharpe ratio to zero-investment strategies and to a fairly general model of returns distributions, where maximizing the Sharpe ratio is not necessarily in the best interest of the fund's investors. By contrast, we assume a positive investment, and we limit the returns distributions under consideration to those for which we can establish a theoretical foundation for the instantaneous Sharpe ratio based on expected utility maximization.

In the case where the value of a fund has constant volatility, we can compare the instantaneous Sharpe ratio with a discrete Sharpe ratio calculated from continuously compounded rates of return. It turns out that the instantaneous Sharpe ratio equals the discrete Sharpe ratio *plus* half of the volatility of the fund.

Unlike the assumption of a constant instantaneous Sharpe ratio, the assumption of a constant volatility rules out various dynamic strategies and option strategies on the part of the fund manager. Moreover, under the joint assumption of constant volatility and a constant instantaneous Sharpe ratio, the excess returns are i.i.d., and the standard estimation theory for the Sharpe ratio derived by Lo (2002) applies. Lo shows that care must be taken in estimating the Sharpe ratio when returns are not i.i.d., such as, for example, when they are autocorrelated.

The relative size of the volatility adjustment to the Sharpe ratio does not depend on whether returns are expressed per day, per month, or per year. The same is true of the ranking of portfolios produced by the instantaneous Sharpe ratio. The fact that the instantaneous Sharpe ratio differs from the discrete Sharpe ratio by half of the volatility of the fund implies that the discrete and instantaneous Sharpe ratios may well produce different rankings of funds. The instantaneous Sharpe ratio does not penalize fund managers as much for taking risks as the discrete ratio does. In particular, if two funds have the same discrete Sharpe ratio but different volatilities, then the fund with higher volatility will be the better performer. The intuition behind this result is that the static one-period theory on which the discrete Sharpe ratio is based overestimates the riskiness of high volatility funds, because it does not take into account the investors' ability to change the fraction of their wealth allocated to the fund over time.

The instantaneous Sharpe ratio does not reward the fund manager for taking risks without regard to the expected rate of return. If he increases the volatility

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

of the fund, then he has to raise the instantaneous excess rate of return at least proportionally to keep the same instantaneous Sharpe ratio, and he has to raise the instantaneous excess return more than proportionally to increase the instantaneous Sharpe ratio. If either the ordinary or the instantaneous Sharpe ratio is used to make statements about whether funds on average have a higher Sharpe ratio than a benchmark, then it is subject to survival bias. This issue pertains equally to ordinary and to instantaneous Sharpe ratios. However, if the Sharpe ratios are used to rank funds, then survival bias should not be an issue.

Jensen's alpha was proposed by Jensen (1968), (1969) and is used by both practitioners and academics. To construct a version of Jensen's alpha that is appropriate in continuous time, we need to interpret it in terms of optimal portfolio choice. If an investor identifies a fund that has a positive alpha, then what exactly does that tell him about how to maximize his expected utility? The literature seems to have been silent on this point, although the following answer is not surprising.

Suppose the investor initially holds a combination of the riskless asset and an index portfolio. He considers whether to tilt his portfolio holdings towards an actively managed fund by investing a small proportion of his wealth in it. He should do so only if it raises his expected utility and, hence, only if it raises the Sharpe ratio of his overall portfolio. We show that Jensen's alpha is proportional to the first derivative of the overall Sharpe ratio with respect to the proportion invested in the active fund. Hence, a positive alpha means that the investor can increase his expected utility by investing at least a small amount in the fund.

This relation between Jensen's alpha and the Sharpe ratio holds in a dynamic model as well as in a static model. In a dynamic model, the relevant version of alpha is the *instantaneous alpha*. It is effectively the same as the discrete alpha, except that in calculating it, the rates of return over finite time intervals are replaced by instantaneous rates of return. We show that the instantaneous alpha is equal to the discrete alpha *plus* half the variance of the portfolio *minus* half the covariance of the portfolio with the benchmark.

The rest of the study is organized as follows. Section II shows how the instantaneous Sharpe ratio can be used as a performance criterion. Section III derives the relation between the instantaneous and the discrete Sharpe ratio. Section IV derives the explicit relation between Jensen's alpha and the Sharpe ratio. Section V discusses the instantaneous Jensen's alpha and its relation to the discrete Jensen's alpha. We conclude in Section VI.

II. The Instantaneous Sharpe Ratio

The discrete Sharpe ratio of a portfolio is the ratio between the expected excess rate of return and the volatility, $S = (Er_p - r_f)/\sqrt{\text{var}(r_p)}$, where r_p is the rate of return on the portfolio, and r_f is the riskless rate. In this section, we define the instantaneous Sharpe ratio and show that it unambiguously ranks funds in a continuous-time setting for expected utility maximizing investors who invest in one fund and in the riskless asset.

For a general introduction to continuous-time finance models, see Nielsen (1999). There is an instantaneously riskless asset with value $M(t) = M(0)$

$\exp\{\int_0^t r\,ds\}$, where r is the instantaneously riskless interest rate. The value process F of a portfolio or fund, with dividends reinvested, is assumed to be a positive Itô process with differential $dF/F = \mu\,dt + \sigma\,dW$, where μ is the instantaneous expected rate of return and σ is the K-dimensional row vector of instantaneous relative dispersion coefficients. The volatility of the fund will be $\sqrt{\sigma\sigma^\top}$. It is assumed to be positive.

The *instantaneous Sharpe ratio* of the fund, denoted by S_{inst}, is defined as $S_{\text{inst}} = (\mu - r)/\sqrt{\sigma\sigma^\top}$. We will assume that the investor chooses one fund and then splits his wealth between this fund and the money market account. The way in which he splits it may change over time in response to new information. In other words, he implements a *portfolio strategy*, which in this context is a one-dimensional process q. The interpretation is that he puts the fraction q of his wealth in the fund and $1 - q$ in the money market account. If $q > 1$, then he uses leverage. The resulting wealth process V has dynamics $dV/V = (q(\mu - r) + r)\,dt + q\sigma\,dW$. We assume a finite time horizon $[0, T]$. The investor chooses q so as to maximize his expected utility of final wealth $V(T)$.

Proposition 1 is the theoretical foundation for using the instantaneous Sharpe ratio for performance measurement in a dynamic framework. It says that an investor who splits his wealth between a money market account and a fund can obtain a higher expected utility the higher is the instantaneous Sharpe ratio of the fund that he chooses, provided that the interest rate varies in a deterministic manner and that the instantaneous Sharpe ratios of the funds under consideration are constant.

It follows that if the investor is choosing one and only one among a number of funds, each of which has a constant instantaneous Sharpe ratio, then he will prefer the one that has the highest instantaneous Sharpe ratio. In this sense, the instantaneous Sharpe ratio can be used to rank funds in the dynamic framework exactly like the discrete Sharpe ratio in a static model.

Proposition 1. Suppose the interest rate r is deterministic. Consider two funds whose price processes F_1 and F_2 have differentials $dF_i/F_i = \mu_i\,dt + \sigma_i\,dW$ for $i = 1, 2$. Suppose the instantaneous Sharpe ratios $S_{\text{inst},i} = (\mu_i - r)/\sqrt{\sigma_i\sigma_i^\top}$, $i = 1, 2$, are positive constants. Given the investor's utility function, the maximum expected utility he can obtain from a portfolio strategy that involves only the fund F_1 and is adapted to $\mathcal{F}$ is strictly larger than the maximum expected utility he can obtain from a portfolio strategy that involves only the fund F_2 and is adapted to $\mathcal{F}$, if and only if $S_{\text{inst},1} > S_{\text{inst},2}$.

The proof of Proposition 1 is in the Appendix. It is not entirely simple for several reasons: i) the Wiener process W is potentially high dimensional, which allows the two funds to be less than perfectly instantaneously correlated, ii) the investor's trading strategy may in principle be contingent on much more information than just observing the value of the fund he is trading, and iii) the relative dispersion vector σ_i of the fund may be stochastically time varying.

Proposition 1 assumes that the instantaneous Sharpe ratios are constant, but it does not assume that the volatilities $\sqrt{\sigma_i\sigma_i^\top}$ or the excess returns $\mu_i - r$ are constant. This has two important implications. First, so long as a fund invests in an underlying portfolio that has constant instantaneous Sharpe ratio, it may well

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

engage in a dynamic strategy with respect to the fraction of asset value invested in the portfolio and the fraction invested in the riskless asset. These fractions may be stochastically time varying and may involve leverage. For example, the fund may follow a strategy of doubling up, or increasing its bets when it suffers losses. Such strategies do not affect the instantaneous Sharpe ratio, so the fund will still have a constant instantaneous Sharpe ratio.

Second, the fund may engage in a strategy that involves buying and selling contingent claims on the underlying portfolio, at least if the latter has constant volatility. If the portfolio has constant volatility, then it conforms to the dynamics underlying the Black-Scholes model, and the excess return and volatility of a contingent claim equals the claim's elasticity times the excess return and volatility, respectively, of the portfolio (see Nielsen (1999), Chapter 6, Section 6.2). Hence, the instantaneous Sharpe ratio of the claim equals the instantaneous Sharpe ratio of the underlying portfolio. The fund will have stochastically time-varying excess return and volatility, but since all the contingent claims are perfectly correlated with the underlying portfolio and have the same constant instantaneous Sharpe ratio, the fund also has that same constant instantaneous Sharpe ratio.

III. Discrete and Instantaneous Sharpe Ratios

In this section, we derive a relation between the instantaneous Sharpe ratio and the discrete Sharpe ratio. We need some manageable assumption about the volatility to calculate the ordinary Sharpe ratio. The simplest assumption that will do is constant volatility. Hence, we assume that the volatility σ and the expected instantaneous excess rate of return $\mu - r$ of the fund are constant. This implies that the instantaneous Sharpe ratio of the fund will be constant. Like in Proposition 1, we assume that the interest rate r is deterministic. Although the expected instantaneous excess return $\mu - r$ is constant, the expected instantaneous return μ itself may not be constant.

The continuously compounded rate of return r_f on the money market account over the time interval $[t, t+\tau]$ is $r_f = \ln M(t+\tau) - \ln M(t) = \int_t^{t+\tau} r\,ds$. It is deterministic. Since we are now considering only one fund at a time, we can assume that the Wiener process is one-dimensional and that σ is one-dimensional.

The mean of the continuously compounded rate of return r_p on the portfolio over the time interval $[t, t+\tau]$ is $Er_p = r_f + m\tau$, where $m = \mu - r - \frac{1}{2}\sigma^2$. The variance is $\sigma^2\tau$, and the standard deviation is $\sigma\sqrt{\tau}$.

The discrete Sharpe ratio is the ratio between the mean and standard deviation of excess rates of return over a discrete period. If the rates of return are expressed as continuously compounded rates per period of length τ, then the discrete Sharpe ratio is $(Er_p - r_f)/\sqrt{\text{var}(r_p)} = (m\tau)/(\sigma\sqrt{\tau}) = S\sqrt{\tau}$, where $S = m/\sigma$ is the discrete Sharpe ratio based on continuously compounded annualized rates.

By substituting the definition of m into the definition of the instantaneous Sharpe ratio, we find the relation between the discrete Sharpe ratio and the instantaneous Sharpe ratio, $S_{\text{inst}} = (\mu - r)/\sigma = (m + \frac{1}{2}\sigma^2)/\sigma = S + \frac{1}{2}\sigma$. So, the instantaneous Sharpe ratio differs from the discrete Sharpe ratio by a bias equal to $\sigma/2$. This bias of course comes from the difference of $\sigma^2/2$ between the instantaneous mean excess return $\mu - r$ and the discrete mean excess return m.

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

It is important to recognize that i) while the discrete and instantaneous Sharpe ratios do depend on whether returns are expressed per day, per month, or per year, the ranking of portfolios that they produce does not, ii) the relative size of the bias does not depend on whether returns are expressed per day, per month, or per year, and iii) when the Sharpe ratios are estimated from data, the importance of the bias is independent of the frequency of the data. To make points i) and ii), we calculate the instantaneous and discrete Sharpe ratios for returns expressed per period of length τ, and then we express the bias as a fraction of the discrete Sharpe ratio.

Observe that the definition of the instantaneous Sharpe ratio as $S_{\text{inst}} = (\mu - r)/\sigma$ is based on instantaneous returns per period of length one, say one year. The instantaneous Sharpe ratio corresponding to rates of return per time period of length τ is $(\mu\tau - r\tau)/(\sigma\sqrt{\tau}) = S_{\text{inst}}\sqrt{\tau} = S\sqrt{\tau} + \frac{1}{2}\sigma\sqrt{\tau}$. It is clear that the rankings of funds produced by $S_{\text{inst}}\sqrt{\tau}$ and $S\sqrt{\tau}$ are independent of τ, which was point i).

The size of the bias is $S_{\text{inst}}\sqrt{\tau} - S\sqrt{\tau} = \frac{1}{2}\sigma\sqrt{\tau}$, which of course goes to zero as the length τ of the time interval goes to zero. However, expressed as a fraction of the discrete Sharpe ratio $S\sqrt{\tau}$, the bias is $(S_{\text{inst}}\sqrt{\tau} - S\sqrt{\tau})/(S\sqrt{\tau}) = (S_{\text{inst}} - S)/S$, which is independent of τ. This was point ii). The relative bias can also be written as $(S_{\text{inst}}\sqrt{\tau} - S\sqrt{\tau})/(S\sqrt{\tau}) = (\mu - r - m)/m = (\mu\tau - r\tau - m\tau)/(m\tau)$. It equals the difference between the instantaneous and the discrete expected excess return per period of length τ, expressed as a fraction of the discrete expected excess return.

Finally, iii) if the Sharpe ratios are estimated from data, then the quality of the estimate will of course be better the more data is used and, in particular, the higher the frequency of the data. However, the true underlying values of the ratios are unaffected, provided that they are expressed in terms of returns per period of a fixed length, such as a year, independently of the sampling frequency.

When estimating the instantaneous Sharpe ratio, we have to take into account the fact that while the parameters μ and σ refer to instantaneous returns, we can actually only observe returns over discrete-time periods such as days, weeks, months, or years. The equation, $S_{\text{inst}} = (m + \frac{1}{2}\sigma^2)/\sigma$, has the virtue of expressing the instantaneous Sharpe ratio in terms of discrete-time moments of the rates of return, since m is the expectation of the annualized discrete-time rate of return and σ is the standard deviation.

The fact that the instantaneous Sharpe ratio equals the discrete Sharpe ratio plus half of the volatility of the fund implies that the ranking of funds based on the discrete and the instantaneous Sharpe ratios may well be different. In particular, if two funds have the same discrete Sharpe ratio but one has higher volatility than the other, then they will be ranked as equal by the discrete Sharpe ratio while the one with higher volatility will be ranked higher by the instantaneous Sharpe ratio. In other words, given the mean annualized excess rate of return m, the instantaneous ratio penalizes the fund manager less than does the discrete ratio for taking risk in the form of volatility. The fund with higher volatility will enable the investor to achieve a higher expected utility.

The intuition behind this result is that the static one-period theory on which the discrete Sharpe ratio is based overestimates the riskiness of high volatility funds, because it does not take into account the investors' ability to change the fraction of their wealth allocated to the fund over time. Take as an example an

investor who wants to hold 50% of his wealth in the fund and 50% in the riskless asset, and whose investment horizon T is one year. In the static framework, he initially invests half of his money in the fund and half in the riskless asset, and then he waits for a year to see what happens. However, already after a month, the value of the fund may have gone up so that he actually holds 60% in the fund and only 40% in the riskless asset. During the course of the year, this situation may be further exacerbated.

By contrast, in the dynamic framework, the investor will immediately react to the increase in the value of the fund by selling some of it and investing the proceeds in the riskless asset, so that he always holds exactly 50% in each. This lowers the overall riskiness of his strategy. The difference is reflected in the modification of the Sharpe ratio.

The fact that the instantaneous Sharpe ratio penalizes the fund manager less for taking risk does not mean that it rewards him for taking risks without regard to the expected rate of return. If he increases the volatility of the fund, then he has to raise the instantaneous excess rate of return $\mu - r$ at least proportionally to keep the same instantaneous Sharpe ratio. To improve its instantaneous Sharpe ratio, and thereby improve its relative ranking, the fund has to increase its instantaneous excess return more than proportionally to any increase in its volatility.

IV. Jensen's Alpha

There are various versions of Jensen's alpha, corresponding to different asset pricing models. Here we will only discuss the original Jensen's alpha, which corresponds to the mean-variance CAPM, and its continuous-time modification. The usual interpretation of alpha is that it is a risk-adjusted performance measure that adjusts expected or average returns for beta risk. However, this interpretation does not explicitly relate alpha to optimal portfolio choice or say precisely what an investor should do if he identifies one or more funds with positive alpha.

This section gives an interpretation of Jensen's alpha in terms of portfolio optimization and explains the relation between Jensen's alpha and the Sharpe ratio. The rates of return in the formulas to follow can be interpreted either as rates of return over discrete-time periods, as will be appropriate in a static model, or as instantaneous rates of return, for use in a dynamic model. The expectations, variances, and covariances should be interpreted accordingly.

Jensen's alpha of a portfolio, relative to an index or benchmark x, is defined as $\alpha = Er_p - r_f - \beta(Er_x - r_f)$, where r_f is the riskless rate, r_p and r_x are the rates of return on the portfolio p and on the index x, and $\beta = \mathrm{cov}(r_p, r_x) / \mathrm{var}(r_x)$ is the beta of the portfolio with respect to the index.

If indeed the index x is efficient, then the true alpha of every security and every portfolio will be zero, although an estimated alpha may be different from zero because of estimation error. However, alpha can be calculated and given a precise interpretation in terms of portfolio optimization even if the index is not efficient.

Suppose the investor initially holds a combination of the riskless asset and an index portfolio tracking the index x, in proportions $1 - \nu$ and ν. He now considers whether to tilt his portfolio a little bit in the direction of the fund p. In other

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

words, he considers taking a small fraction ϵ of his wealth and investing it in the portfolio p, while reducing the fractions held in the riskless asset and the index to $(1-\epsilon)(1-\nu)$ and $(1-\epsilon)\nu$, respectively.

Let $S(\epsilon)$ denote the Sharpe ratio (or instantaneous Sharpe ratio) of the new portfolio.

Proposition 2. The derivative of $S(\epsilon)$ with respect to ϵ, evaluated at $\epsilon = 0$, is $S'(0) = \alpha/(\nu\sqrt{\text{var}(r_x)})$.

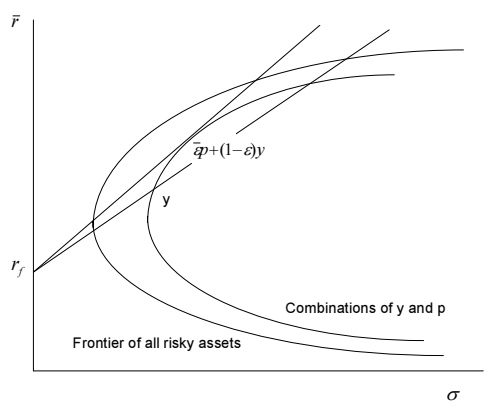
The proof of Proposition 2 is in the Appendix.

Proposition 2 leads to the following interpretation of alpha. If $\alpha > 0$, then an investor who basically invests in the index or in a combination of the index and the riskless asset can increase his Sharpe ratio and hence his expected utility by investing a small positive amount in the fund p. Of course, if $\alpha < 0$, then he can achieve the same effect by short-selling the fund, if this is possible.

When ϵ varies, the standard deviation and mean of the investor's entire portfolio traces out a hyperbolic curve, which is in fact the risky portfolio frontier generated by two assets, the initial portfolio and the fund p. We illustrate this frontier in Figure 1, where $\alpha > 0$. The frontier should not be confused with the usual frontier constructed from all available securities.

FIGURE 1

The Frontier Generated by the Index and the Fund



When $\epsilon = 0$, we are at the point y. As ϵ increases, we move up along the upper branch of the small hyperbola. The Sharpe ratio $S(\epsilon)$ is initially increasing and then decreasing. It has a maximum point $\bar{\epsilon} > 0$, which represents the optimal fraction of wealth to take out of the initial portfolio and put into the fund p. It corresponds to the point $\bar{\epsilon}p + (1-\bar{\epsilon})y$ in Figure 1.

It is alternatively possible that $S(\epsilon)$ does not have a maximum but is increasing for all $\epsilon > 0$. This occurs if the riskless rate is at or above the expected rate of return on the minimum variance portfolio formed from the index and the fund p, which corresponds to the top-point of the small hyperbolic curve in the figure. This resembles the situation where the riskless rate is at or above the global

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

minimum variance portfolio formed from the risky securities, as analyzed, for example, in Huang and Litzenberger (1988).

Observe that alpha alone does not say how much the investor should optimally invest in the fund. In other words, we cannot calculate $\bar{\epsilon}$, the optimal value of ϵ, knowing only the value of alpha. The idiosyncratic variance of the fund also matters. If the investor puts too large a fraction of his wealth into the fund p, then the idiosyncratic risk may result in a lower Sharpe ratio and a lower expected utility.

The analysis above applies not only in a static model but also in a continuous-time model, when the rates of return over a discrete-time interval are replaced by instantaneous rates of return. The Sharpe ratio will be replaced by an instantaneous Sharpe ratio, and alpha will be replaced by an instantaneous alpha, which we shall define in the following section.

V. The Instantaneous Alpha

In this section, we define the instantaneous alpha and derive a relation between the instantaneous and the discrete alphas. Let F_x be the value of the index fund with dividends reinvested, and let F_p be the value of the other fund with dividends reinvested. Assume that they follow the processes $F_x(t) = F_x(0)\exp\{\int_0^t(\mu_x - \frac{1}{2}\sigma_x^2)ds + \int_0^t \sigma_x\, dZ_x$ and $F_p(t) = F_p(0)\exp\{\int_0^t(\mu_p - \frac{1}{2}\sigma_p^2)ds + \int_0^t \sigma_p\, dZ_p\}$, where μ_x and μ_p are the instantaneous expected rates of return, σ_x and σ_p are the instantaneous volatilities or standard deviations of the rates of return, and Z_1 and Z_2 are two potentially correlated standard Wiener processes with correlation coefficient ρ.

The *instantaneous alpha* of the fund, denoted by α_{inst}, is defined as $\alpha_{\text{inst}} = \mu_p - r - \beta_{\text{inst}}(\mu_x - r)$, where $\beta_{\text{inst}} = (\sigma_p\sigma_x\rho)/\sigma_x^2 = (\sigma_p\rho)/\sigma_x$. The instantaneous alpha is effectively the discrete alpha with the rates of return over finite time intervals replaced by instantaneous rates of return.

For the purpose of deriving a relation between the instantaneous and the discrete alpha, assume that the interest rate varies in a deterministic manner, and that the correlation ρ, the volatilities σ_p and σ_x, and the instantaneous expected excess rates of return $\mu_p - r$ and $\mu_x - r$ are constant.

The means of the continuously compounded rates of return r_x and r_p on x and p over the time interval $[t, t+\tau]$ are $Er_p = r_f + m_p\tau$ and $Er_x = r_f + m_x\tau$, where $m_x = \mu_x - r - \frac{1}{2}\sigma_x^2$ and $m_p = \mu_p - r - \frac{1}{2}\sigma_p^2$. The variances and the covariance are $\text{var}(r_p) = \sigma_p^2\tau$, $\text{var}(r_x) = \sigma_x^2\tau$, and $\text{cov}(r_p, r_x) = \sigma_p\sigma_x\rho\tau$. If the rates of return are expressed per period of length τ, then the discrete Jensen's alpha is $m_p\tau - ((\text{cov}(r_p, r_x))/(\text{var}(r_x)))m_x\tau = m_p\tau - ((\sigma_p\sigma_x\rho\tau)/(\sigma_x^2\tau))m_x\tau = (m_p - \beta m_x)\tau = \alpha\tau$, where $\beta = \text{cov}(r_p, r_x)/\text{var}(r_x) = (\sigma_p\sigma_x\rho\tau)/(\sigma_x^2\tau) = (\sigma_p\sigma_x\rho)/\sigma_x^2 = \beta_{\text{inst}}$ and $\alpha = m_p - \beta m_x$ is the discrete Jensen's alpha based on annualized returns.

By substituting the definitions of m_p and m_x into the definition of the instantaneous alpha, we find the relation between the discrete alpha and the instantaneous alpha, $\alpha_{\text{inst}} = \alpha + \frac{1}{2}(\sigma_p^2 - \sigma_p\sigma_x\rho)$. It follows that the instantaneous Jensen's alpha differs from the discrete Jensen's alpha by a bias equal to $(\sigma_p^2 - \sigma_p\sigma_x\rho)/2$.

Similar to the Sharpe ratios, it is important to recognize that i) while the discrete and instantaneous alphas do depend on whether returns are expressed per

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

day, per month, or per year, the ranking of portfolios that they produce does not, ii) the relative size of the bias does not depend on whether returns are expressed per day, per month, or per year, and iii) if the alphas are estimated from data, then the importance of the bias is independent of the frequency of the data.

To make points i) and ii), we calculate the instantaneous alpha for returns expressed per period of length τ, and then we express the bias as a fraction of the discrete alpha. Observe that the definition of the instantaneous alpha as $\alpha_{\text{inst}} = \mu_p - r - \beta_{\text{inst}}(\mu_x - r)$ is based on instantaneous returns per period of length one, say one year. The instantaneous alpha corresponding to rates of return per time period of length τ is $(\mu_p - r)\tau - \beta_{\text{inst}}(\mu_x - r)\tau = \alpha_{\text{inst}}\tau$. It is clear that the rankings of funds produced by $\alpha_{\text{inst}}\tau$ and $\alpha\tau$ are independent of τ. This illustrates point ii). The size of the bias is $\alpha_{\text{inst}}\tau - \alpha\tau = \frac{1}{2}(\sigma_p^2 - \sigma_p\sigma_x\rho)\tau$, which goes to zero as the length τ of the time interval goes to zero. However, expressed as a fraction of the discrete alpha, the bias is $(\alpha_{\text{inst}}\tau - \alpha\tau)/(\alpha\tau) = (\alpha_{\text{inst}} - \alpha)/\alpha$, which is independent of τ. This demonstrates point ii). Finally, iii) if the alphas are estimated from data, then the same arguments made for the instantaneous and discrete Sharpe ratios apply.

Similar to the case of the instantaneous Sharpe ratio, the equation, $\alpha_{\text{inst}} = \alpha + \frac{1}{2}(\sigma_p^2 - \sigma_p\sigma_x\rho)$, expresses the instantaneous alpha in terms of discrete-time moments of the rates of return. This is useful when estimating it from data.

VI. Conclusions

This paper has proposed modifications of the Sharpe ratio and Jensen's alpha, which are consistent with expected utility maximization in a continuous-time model. Specifically, the modifications take into account the fact that investors may change the split of their wealth between the fund and the riskless asset over time. The theory assumes that the Sharpe ratios are constant, but it allows for stochastically time-varying volatilities, which could arise, for example, from a dynamic leverage strategy or from a strategy of buying contingent claims such as puts and calls on an underlying portfolio.

The instantaneous Sharpe ratio does not necessarily deliver the same ranking of funds as its discrete version. In fact, in the special case of constant volatility, we related these two versions and found that the instantaneous Sharpe ratio penalizes a fund less for taking risk than does the discrete ratio.

We derived a precise interpretation of Jensen's alpha in terms of optimal portfolio choice by relating it to the Sharpe ratio. Specifically, a positive alpha of a fund means that an investor who initially holds a benchmark index fund can improve his Sharpe ratio by diverting a small fraction of his wealth into the fund.

The modified performance evaluation criteria proposed in this paper have been derived under the simplest possible assumptions. There is scope to explore the modifications to the theory required when the Sharpe ratios change over time. There is also scope to explore the estimation of all the performance measures, in a way that would be consistent with theory, when the volatilities and expected excess returns of the funds are not constant. Such extensions go beyond the boundaries of this paper.

VII. Appendix: Formalities and Proofs

The investors' information structure is represented by a filtration $(\mathcal{F}_t)_{t\in[0,T]}$ on an underlying probability space $(\Omega, \mathcal{F}, P)$. The interpretation is that $\mathcal{F}_t$ is the information set available to the investors at time t. Random fluctuations in securities prices are driven by a K-dimensional process W, which is a K-dimensional Wiener process with respect to the filtration. Portfolio strategies, and the instantaneous means and dispersions of value processes, are assumed to be measurable and adapted to the filtration.

The proof of Proposition 1 relies on the following lemma.

Lemma 1. Suppose the interest rate r is deterministic. Let B be a one-dimensional standard Brownian motion, and let $\mathcal{F}^B$ be the augmented filtration generated by B. Consider two funds whose price processes F and $\hat{F}$ have differentials $(dF)/F = \mu\,dt + \sigma\,dW$ and $(d\hat{F})/\hat{F} = (s^2 + r)\,dt + s\,dB$, where $s = (\mu - r)/\sqrt{\sigma\sigma^\top}$ is assumed to be constant. Given the investor's utility function, the maximum expected utility he can obtain from a portfolio strategy that involves only fund F and is adapted to $\mathcal{F}$ is the same as the maximum expected utility he can obtain from a portfolio strategy that involves only fund $\hat{F}$ and is adapted to $\mathcal{F}^B$.

Proof. Set $\phi = s/\sqrt{\sigma\sigma^\top} = (\mu - r)/(\sigma\sigma^\top)$ and $\lambda = \phi\sigma$, then $\lambda\lambda^\top = s^2$. Define a one-dimensional standard Brownian motion C by $C(t) = \int_0^t (1/s)\lambda\,dW$ and let $\mathcal{F}^C$ be the augmented filtration generated by C. It follows from the results in Nielsen and Vassalou (1997) that the optimal portfolio strategy when trading fund F has the form $q = a\phi$, where a is a process that is adapted to $\mathcal{F}^C$. A strategy of this form gives the following dynamics of wealth, $(dV)/V = (q(\mu - r) + r)\,dt + q\sigma\,dW = (a\phi s\sqrt{\sigma\sigma^\top} + r)\,dt + a\phi\sigma\,dW = (as^2 + r)\,dt + a\lambda\,dW = (as^2 + r)\,dt + as\,dC$. Consider fund F^ϕ, which arises from trading fund F using portfolio strategy ϕ. It has dynamics $(dF^\phi)/F^\phi = (s^2 + r)\,dt + s\,dC$. The wealth dynamics arising from trading fund F using portfolio strategy $a\phi$ is $(dV)/V = (as^2 + r)\,dt + as\,dC$, which is the same as the wealth dynamics arising from trading fund F^ϕ using portfolio strategy a. Hence, the maximum expected utility from trading fund F using portfolio strategies that are adapted to $\mathcal{F}$ is identical to the maximum expected utility from trading fund F^ϕ using portfolio strategies that are adapted to $\mathcal{F}^C$. The latter is obviously identical to the maximum expected utility from trading fund $\hat{F}$ using portfolio strategies that are adapted to $\mathcal{F}^B$. □

Proof of Proposition 1. Let B be a one-dimensional standard Brownian motion, and let $\mathcal{F}^B$ be the augmented filtration generated by B. For each s, consider a fund whose price processes $\hat{F}[s]$ have differential $(d\hat{F}[s])/\hat{F}[s] = (s^2 + r)\,dt + s\,dB$. According to Lemma 1, given the investor's utility function, the maximum expected utility he can obtain from a portfolio strategy that involves only fund F_i and is adapted to $\mathcal{F}$ is the same as the maximum expected utility he can obtain from a portfolio strategy that involves only fund $\hat{F}[S_{\text{inst},i}]$ and is adapted to $\mathcal{F}^B$. Therefore, what we need to show is that if $S_{\text{inst},1} > S_{\text{inst},2}$, then the maximum expected utility the investor can obtain from trading in fund $\hat{F}[S_{\text{inst},1}]$ with a portfolio strategy that is adapted to $\mathcal{F}^B$ is strictly larger than the maximum expected utility he can obtain from trading in fund $\hat{F}[S_{\text{inst},2}]$ with a portfolio strategy that is adapted to $\mathcal{F}^B$.

Let a be the optimal portfolio strategy when he trades fund $\hat{F}[S_{\text{inst},2}]$. Since this fund has constant instantaneous mean and dispersion, it is known from Merton (1971) that a has the form, $a = \gamma(\mu_2 - r)/(\sigma_2\sigma_2^\top) = \gamma(S_{\text{inst},2}/\sqrt{\sigma\sigma^\top})$, where $\gamma > 0$ is the relative risk tolerance of the investor's utility function. Hence, $a > 0$. If the investor uses the portfolio strategy a to trade the fund $\hat{F}[S_{\text{inst},i}]$, then his wealth dynamics is $(dV_i)/V_i = (aS^2_{\text{inst},i} + r)\,dt + aS_{\text{inst},i}\,dB$. The logarithm of his final wealth will be $\ln V_i(T) = \ln V(0) + \int_0^t (aS^2_{\text{inst},i} + r - \tfrac{1}{2}a^2)\,dt + \int_0^t aS_{\text{inst},i}\,dB$. Hence, $\ln V_1(T) - \ln V_2(T) = \int_0^t a(S^1_{\text{inst},i} - S^2_{\text{inst},i})\,dt > 0$. Therefore, the investor obtains a strictly higher expected utility by using a to trade fund $\hat{F}[S_{\text{inst},1}]$ than by using it to trade the fund $\hat{F}[S_{\text{inst},2}]$. □

Proof of Proposition 2. Let y denote the initial portfolio. Its rate of return is $r_y = (1 - \nu)r_f + \nu r_x$. The expected excess rate of return and variance of the new portfolio will be

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press

$Er_y - r_f + \epsilon(Er_p - Er_y)$ and $(1-\epsilon)^2 \operatorname{var}(r_y) + \epsilon^2 \operatorname{var}(r_p) + 2\epsilon(1-\epsilon)\operatorname{cov}(r_y, r_p)$, respectively, and the Sharpe ratio (or instantaneous Sharpe ratio) will be

$$S(\epsilon) = \frac{Er_y - r_f + \epsilon(Er_p - Er_y)}{\sqrt{(1-\epsilon)^2 \operatorname{var}(r_y) + \epsilon^2 \operatorname{var}(r_p) + 2\epsilon(1-\epsilon)\operatorname{cov}(r_y, r_p)}}.$$

Set $E(\epsilon) = Er_y - r_f + \epsilon(Er_p - Er_y)$, $v(\epsilon) = (1-\epsilon)^2 \operatorname{var}(r_y) + \epsilon^2 \operatorname{var}(r_p) + 2\epsilon(1-\epsilon)\operatorname{cov}(r_y, r_p)$, and $\sigma(\epsilon) = \sqrt{v(\epsilon)}$, Then

$$S(\epsilon) = \frac{Er_y - r_f + \epsilon(Er_p - Er_y)}{\sqrt{(1-\epsilon)^2 \operatorname{var}(r_y) + \epsilon^2 \operatorname{var}(r_p) + 2\epsilon(1-\epsilon)\operatorname{cov}(r_y, r_p)}} = \frac{E(\epsilon)}{\sigma(\epsilon)}.$$

To calculate $S'(0)$, first observe the following, $E(0) = Er_p - Er_y$, $v(0) = \operatorname{var}(r_y)$, $\sigma(0) = \sqrt{\operatorname{var}(r_y)}$, $E'(\epsilon) = Er_p - Er_y$, $v'(\epsilon) = -2(1-\epsilon)\operatorname{var}(r_y) + 2\epsilon \operatorname{var}(r_p) + 2(1-2\epsilon)\operatorname{cov}(r_y, r_p)$, $v'(0) = -2\operatorname{var}(r_y) + 2\operatorname{cov}(r_y, r_p)$, $\sigma'(\epsilon) = \frac{1}{2}(v'(\epsilon))/(\sigma(\epsilon))$, and

$$\sigma'(0) = \frac{1}{2}\frac{v'(0)}{\sigma(0)} = \frac{1}{2}\frac{-2\operatorname{var}(r_y) + 2\operatorname{cov}(r_y, r_p)}{\sqrt{\operatorname{var}(r_y)}} = \frac{-\operatorname{var}(r_y) + \operatorname{cov}(r_y, r_p)}{\sqrt{\operatorname{var}(r_y)}}.$$

Moreover, $Er_p - r_f + (\operatorname{cov}(r_p, r_y)/\operatorname{var}(r_y))(Er_y - r_f) = Er_p - r_f + (\nu \operatorname{cov}(r_p, r_x)/\nu^2 \operatorname{var}(r_x))(\nu Er_x + (1-\nu)r_f - r_f) = Er_p - r_f + (1/\nu)\beta\nu(Er_x - r_f) = Er_p - r_f + \beta(Er_x - r_f) = \alpha$. Now,

$$\begin{aligned} S'(0) &= \frac{E'(0)\sigma(0) - E(0)\sigma'(0)}{v(0)} \\ &= \frac{1}{\operatorname{var}(r_y)}\left[(Er_p - Er_y)\sqrt{\operatorname{var}(r_y)} - (Er_y - r_f)\frac{-\operatorname{var}(r_y) + \operatorname{cov}(r_p, r_y)}{\sqrt{\operatorname{var}(r_y)}}\right] \\ &= \frac{1}{\sqrt{\operatorname{var}(r_y)}}\left[Er_p - Er_y + (Er_y - r_f) - (Er_y - r_f)\frac{\operatorname{cov}(r_p, r_y)}{\operatorname{var}(r_y)}\right] \\ &= \frac{1}{\nu\sqrt{\operatorname{var}(r_x)}}[Er_p - r_f + \beta(Er_y - r_f)] = \frac{\alpha}{\nu\sqrt{\operatorname{var}(r_x)}}. \quad \square \end{aligned}$$

References

Goetzman, W.; J. E. Ingersoll; M. I. Spiegel; and I. Welch. "Sharpening Sharpe Ratios." Working Paper 02-08, Yale IFC (February 2002).

Huang, C., and R. Litzenberger. *Foundations for Financial Economics*. Amsterdam: North-Holland (1988).

Jensen, M. C. "The Performance of Mutual Funds in the Period 1945–1964." *Journal of Finance*, 23 (1968), 389–415.

_____. "Risk, the Pricing of Capital Assets, and the Evaluation of Investment Portfolios." *Journal of Business*, 42 (1969), 167–242.

Lo, A. W. "The Statistics of Sharpe Ratios." *Financial Analysts Journal*, 58 (2002), 36–52.

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

Nielsen, L. T. *Pricing and Hedging of Derivative Securities*. Oxford, U.K.: Oxford Univ. Press (1999).

Nielsen, L. T., and M. Vassalou. "Portfolio Selection with Randomly Time-Varying Moments: The Role of the Instantaneous Capital Market Line." http://www.gsb.columbia.edu/faculty/mvassalou/. Working Paper, Columbia Univ. (1997).

Sharpe, W. F. "Mutual Fund Performance." *Journal of Business*, 34 (1966), 119–138.

https://doi.org/10.1017/S0022109000003902 Published online by Cambridge University Press