

Anat R. Admati
Paul Pfleiderer

Stanford University

Does It All Add Up? Benchmarks and the Compensation of Active Portfolio Managers*

I. Introduction

Suppose a portfolio manager investing in Japanese stocks realizes a return of 27%. How well has she performed? A good answer to this question cannot be given without knowing how well Japanese stocks performed in the aggregate. If we view the problem of assessing this manager's performance as a pure inference problem—after the fact one simply wants to know whether the manager has demonstrated skill in selecting Japanese stocks—then it clearly seems more sensible to consider the difference between the manager's return and that earned on some appropriately chosen Japanese benchmark portfolio than to consider only the unadjusted level of the manager's return.

Since netting out a benchmark portfolio's return from the return of a managed portfolio seems useful in the pure inference problem, it might be concluded that the compensation of portfolio managers should be based on similar adjustments to the total returns they realize. Indeed, there has

* We would like to thank Michael Brennan, Lee Nelson, Bill Sharpe, Doug Diamond (the editor), an anonymous referee and seminar participants at the University of California, Los Angeles, for helpful comments. We are also grateful to the National Science Foundation (grant no. SBR-9308238) and the Financial Research Initiative at the Stanford Graduate School of Business for financial support.

In this article we examine theoretically the use of benchmark portfolios in the compensation of privately informed portfolio managers. We find that the use of a benchmark, and particularly the types of benchmarks often observed in practice, cannot be easily rationalized. Specifically, commonly used benchmark-adjusted compensation schemes are generally inconsistent with optimal risk-sharing and do not lead to the choice of an optimal portfolio for the investor. Moreover, benchmarks do not help in solving potential contracting problems such as inducing the manager to expend effort or trying to screen out uninformed managers.

been a substantial move toward making the compensation of portfolio managers explicitly depend on their performance relative to benchmarks, that is, to use what we will call “benchmark-adjusted compensation.” Even if no explicit adjustment is made, the compensation of portfolio managers often depends implicitly on their performance relative to a benchmark in the sense that the amount of money under their management is affected by such comparisons.

It is generally taken for granted (and perhaps even seems obvious) that benchmark-adjusted compensation is a good idea, and much of the recent discussion on managers’ compensation is concerned with how the benchmark portfolios should be chosen.¹ Less attention has been focused on answering two important and prior questions: do benchmark-adjusted compensation schemes lead to better outcomes for investors than other forms of compensation, and, if so, how do the gains come about?

In this article we examine theoretically the effects of benchmark-adjusted compensation. Specifically, we examine whether a portfolio manager can be induced to choose the optimal portfolio for the investor through the use of benchmarks and whether benchmarks might help solve various types of contracting problems that potentially exist when an investor delegates investment decisions to a portfolio manager. Based on our analysis, we find strong grounds to question the proposition that benchmark-adjusted compensation schemes are useful. In particular, we find that commonly used benchmark-adjusted compensation schemes

- (i) are generally inconsistent with optimal risk sharing;
- (ii) are generally inconsistent with the goal of obtaining the optimal portfolio for the investor, and particularly in coordinating the manager’s portfolio choices with other portfolio holdings of the investor;
- (iii) tend to weaken a manager’s incentives to expend effort (or are at best irrelevant in this context);
- (iv) are not useful in screening out bad (i.e., uninformed) managers; and
- (v) play no role in aligning the manager’s preferences with the investor’s when constraints are placed on the manager’s portfolio choice due to uncertainty about the manager’s risk tolerance.

In other words, in all the contexts we examine in which benchmarks might be valuable, we find that the use of a benchmark, and particularly the types of benchmarks often observed in practice, cannot be easily

1. See, e.g., Bailey (1990); Rennie and Cowhey (1990); Bailey and Tierney (1993); and Gastineau (1994).

rationalized. This is not to say that the notion of “performance-based” fees is flawed. Making compensation depend on the total return realization on the managed portfolio is potentially valuable for aligning the manager’s preferences with those of the investor and for motivating her to expend costly effort to collect information. It is the use of benchmark adjustments in compensation that we question.

Our analysis is conducted in a model (introduced in Sec. II) where the portfolio manager is “active” in the sense that she has some useful private information about securities’ returns. In most of our analysis we consider the case in which this private information concerns the idiosyncratic or firm-specific part of stocks’ returns. In other words, we assume that the manager has valuable information for selecting stocks but no information about the pervasive or market-wide determinants of stock price movements. This is a natural setting for considering the value of using benchmarks because it is this idiosyncratic type of information that is supposedly generating excess returns over and above those of the prespecified benchmark.

In an insightful paper Roll (1992) showed that, in a setting in which the manager has no private information, the portfolio choice she makes while responding to a benchmark-adjusted compensation structure will not be mean-variance efficient.² In Roll’s analysis there is no distinction between the mean-variance frontier (of unadjusted returns) assessed by the manager and that assessed by the investor or sponsor. In our model, by contrast, the manager has private information, and so the manager perceives a different efficient frontier than uninformed investors. Ideally for the investor, the manager will choose the optimal portfolio for the investor on the *conditional* mean-variance frontier (conditional on the particular realization of the manager’s information). This optimal portfolio obviously varies with the information realization and with probability equal to one is not on the *unconditional* mean variance frontier (the one based on the unconditional distribution of returns).³

We start our analysis (in Sec. III) by extending Roll’s result to the case where private information is explicitly included in the model. We show that the use of market-type benchmarks in compensation can lead to a substantial loss in the value of a manager’s information to the investor. For example, in a setting where a manager has information worth about 6% per annum (i.e., if used optimally the information produces a risk-adjusted expected return 6% higher than the risk-adjusted expected return an uninformed investor could attain), we find that the benchmark-adjusted compensation (where the benchmark is equal to the passive portfolio that an uninformed investor would hold) reduces

2. This is true as long as the benchmark is not itself mean-variance efficient.

3. For related discussions on the differences between conditional and unconditional mean variance analysis see, e.g., Admati and Ross (1985), and Hansen and Richard (1987).

the gain by 60%—instead of getting a risk-adjusted expected return 6% higher, the investor only realizes a gain of 2.4%. For less precise information the results are even more striking, and we find that the loss due to the use of such a benchmark may even exceed 100%. In such cases the investor is worse off using the informed manager than investing passively based on no private information. At the same time we show that full value of the manager's information can in fact be realized if no benchmark is used, that is, if compensation is based only on the total unadjusted returns on the manager's portfolio.

Benchmarks are often used by large sponsors whose portfolios are managed by multiple managers. In Section IV we examine the case in which the investor holds part of his wealth in exogenously specified passive portfolio holdings and wants to coordinate the active portfolio manager's portfolio decisions with the passive portfolio holdings. We show that a benchmark can be chosen to "offset" any passive portfolio and lead to as overall optimal choice for the investor, but the benchmarks that are typically used in practice are inconsistent with this type of coordination and again lead to losses in the value of the manager's information.

In constructing a compensation scheme for the manager, the investor is presumably concerned not only with the resulting portfolio choices made by the manager but also with the cost to him of compensating the manager. It might be thought that the investor can lower this cost by shielding the manager from bearing the risk of the benchmark return and that this can be done through the use of benchmark-adjusted compensation. However, in Section V we argue that in general introducing a benchmark adjustment in the manager's compensation interferes with efficient risk sharing and actually increases the cost of compensation.

One common intuition regarding the value of using benchmarks in compensation comes from the agency literature, where it is argued that in creating incentives for managers to expend effort one should "net out" elements that are not under the manager's control. In Section VI we discuss the role of benchmark portfolios in the context of the effort incentive problem, that is, motivating the manager to expend effort to collect the appropriate information. We show that benchmarks have no positive role to play in sharpening incentives. For the most part the presence of a benchmark in the compensation and especially the precise composition of the benchmark is irrelevant for the effort incentive problem. Moreover, we find that not only do benchmark-adjusted compensation plans not strengthen incentives, there are many cases in which they actually weaken them.

Having motivated the use of benchmarks in the context of the pure inference problem of assessing the manager's skill *ex post*, we examine (in Sec. VII) their usefulness as part of the compensation scheme when the manager's skill is unknown to the investor. We first show that

benchmark-adjusted returns are not a sufficient statistic for assessing the quality of a manager's information. We then show that benchmark-adjusted compensation does not enhance the ability of the investor to make the assessment of the manager's skill. Again, the presence and composition of the benchmark are largely irrelevant to the ability to make the inference. Finally, we show that use of a benchmark in the explicit compensation of the manager can help counteract the effects of benchmark-adjusted compensation that is implicit, that is, managers who outperform their benchmarks are more likely to get more funds under management than those who do not. However, the explicit benchmarks that should be used to accomplish this are dramatically different from those used in practice.

Portfolio managers often face explicit constraints on the riskiness of their portfolios. For example, as discussed by Roll (1992), the investor may require that the portfolio have a beta equal to some given level. In Section VIII we examine the role of benchmarks when such constraints are imposed. In particular we examine the case where constraints are imposed because the investor is uncertain about the manager's risk tolerance. We show that the use of a benchmark does not make the portfolio choice better for the investor, and it generally makes it worse when constraints are present.

In addition to Roll (1992), a number of papers discuss issues related to our findings in this paper. Sharpe (1981) addresses the problem of coordinating investments when different parts of a portfolio are managed by different managers. He allows the investor to alter the objective functions of managers to try to achieve optimality. We will make most of our points in the context where there is just one manager, but the observations in Sharpe (1981) are similar to some that we make both in the context of inducing optimal portfolio choices and in the context of specialization. Brennan (1993) assumes that the manager's compensation is based on performance relative to an index and examines the equilibrium that results when a substantial part of the market is invested by intermediaries responding to such contracts. Our analysis in this article is a partial equilibrium analysis that looks more closely at the compensation structure. Obviously, taking the analysis to an equilibrium level, as is done in Brennan (1993), is an important next step. Finally, Huddart (1994) and Heinkel and Stoughton (1994) examine dynamic portfolio management contracts, each in a somewhat different model. Both focus on the screening problem and the notion of reputation.

II. The Basic Model

The economy we analyze has n risky securities that are traded in the first period and pay off in the second. Some of our results depend on

whether or not a riskless security exists, and our model allows us to consider both cases. If there is a riskless security, we will normalize its return in much of our analysis to be zero without loss of generality. Absent any private information, all agents assess that the returns of the risky securities, given by the vector \tilde{r} , are jointly normally distributed with an *unconditional* expected return vector μ_u and an *unconditional* variance-covariance matrix V_u .⁴ In some of our analysis we will assume that there is a factor model that generates returns, that is, that returns are given by

$$\tilde{r} = \mu_u + B\tilde{f} + \tilde{\epsilon}, \quad (1)$$

where B is an $n \times k$ matrix of constants, \tilde{f} is a k dimensional vector, and $\tilde{\epsilon}$ and \tilde{f} are independent and have a joint normal distribution with zero mean.

Investment decisions in our model will be made on behalf of an investor who does not observe any private information. This investor delegates the investment decisions to an “active” portfolio manager. (As discussed in the concluding remarks, our results have a number of immediate implications for the more complicated case of multiple managers.) The manager bases her investment decisions on private information she receives concerning the returns of the risky securities. Specifically, the manager observes the signal \tilde{s} , which has a joint normal distribution with the return vector \tilde{r} .

Under the joint normality assumption, the conditional distribution of returns given the manager’s private information is also normal, with a *conditional* expected return vector $\tilde{\mu}_c = E(\tilde{r}|\tilde{s})$ and a *conditional* variance-covariance matrix V_c . Note that V_c does not depend on the observation of \tilde{s} but only on the parameters of the joint distribution of \tilde{r} and \tilde{s} . The conditional expected return vector $\tilde{\mu}_c$ depends on \tilde{s} , however. If the manager observes information only about the idiosyncratic terms in a factor model of returns, then $E(\tilde{f}|\tilde{s}) = 0$ and $E(\tilde{\epsilon}|\tilde{s}) = \tilde{\eta}_c$, where $\tilde{\eta}_c$ has an unconditional expected return of zero.

We normalize the investor’s total wealth to be equal to one and assume that he has constant absolute risk aversion with risk tolerance parameter (inverse of the coefficient of risk aversion) equal to τ . In examining the effect of various compensation structures on the investor’s expected utility we will initially ignore the cost to the investor of compensating the manager and focus on the effect of the manager’s portfolio choices on the resulting wealth distribution of the investor.

4. The expected return vector is given exogenously. That is, we do not consider the effect of the investments made by the investors in our model on equilibrium return distributions. This partial equilibrium analysis is justified if the agents in our model are very small relative to the entire market. While beyond the scope of this article, the equilibrium implications of delegated portfolio management are an important topic for future research. For an important step in this direction see Brennan (1993).

In Sections V and VI we look at the cost of compensating the manager in more detail.

We assume that the manager's preferences are also summarized by a negative exponential utility function over final wealth with a risk tolerance coefficient denoted by ρ . The manager's final wealth depends on the compensation she receives from the investor for her portfolio management activities.⁵ For most of our analysis we assume that the compensation of the manager is of the general form

$$h_1 x' \tilde{r} + h_2 (x - b)' \tilde{r}, \quad (2)$$

where x is the portfolio of risky securities chosen by the manager, b is a benchmark portfolio of risky securities, whose weights are specified in advance, and h_1 and h_2 are constants also specified in advance. The weights in portfolios x and b sum to one if there is no riskless security and are unrestricted if there is a riskless security. Note that the compensation includes two components. First, there is a component that depends on the total return on the manager's portfolio of risky securities $x' \tilde{r}$. Second, there is a benchmark-based component, where the manager receives a fraction of the difference between her portfolio's return and the return on a prespecified benchmark portfolio b . For simplicity we do not include a constant component in the compensation function, but none of our results would change if the compensation also included a fixed payment that does not depend on the returns of the risky securities. The coefficients h_1 and h_2 , as well as the benchmark b , can be customized in the sense that they can depend, for example, on the manager's expertise or style, the precision of her information, her risk tolerance coefficient and that of the investor, and so on. We further assume that $h_1 + h_2 \neq 0$.⁶ Note that this compensation function is linear in the returns on the manager's portfolio.⁷

In order to focus on the issue of coordination, we will initially abstract away from any agency problems by assuming that the investor knows the distributional properties of the information of the manager, as well as her preference parameter. Sections VI and VII discuss some

5. We ignore other sources of income in the analysis, but our basic results should not change if endowment wealth is included for the manager. The case in which each manager might work for multiple investors of the type we consider is more complicated and obviously important, but it is beyond the scope of this article.

6. If $h_1 + h_2 = 0$, then the manager's compensation does not depend on the returns on the risky securities, and therefore her portfolio choices will not be affected by the compensation. This is clearly unrealistic.

7. For example, we do not consider the option effect of compensation schedules that only provide compensation for positive returns relative to the benchmark. See Grinblatt and Titman (1989) for a discussion of these types of contracts and their effect on portfolio choices. It is also possible that the compensation structure penalizes the manager severely when she underperforms relative to the benchmark, while it is not as sensitive to the degree to which she may overperform the benchmark. This, too, would introduce nonlinearities into the compensation structure.

issues that arise when the precision of the manager's information is either determined by the manager's effort or is unknown to the investor. Section VIII addresses the case where the manager's risk tolerance might not be known to the investor.

III. The Effect of Benchmark-Adjusted Compensation

We now show that the use of a benchmark portfolio in the compensation of the portfolio manager generally leads to a suboptimal portfolio choice for the investor and therefore to a loss in the value of the manager's information to the investor. We show that the economic magnitude of this loss can be quite substantial in realistic examples. These results are related to the analysis of Roll (1992), but here we explicitly model the manager as having private information, and so we work with the conditional distributions of returns, which are different from the unconditional distributions. The effects of any compensation structure are then measured by how they affect the ex ante expected utility of the investor.

We start by noting that in our setting, if the investor could observe the private information of the manager, his optimal portfolio conditional on this information, which we denote by \tilde{x}_c^* , would be given by

$$\tilde{x}_c^* = \theta_c + \tau \Delta_c \tilde{\mu}_c, \quad (3)$$

where θ_c is the global minimum variance portfolio obtained using the *conditional* variance-covariance matrix V_c , and Δ_c is a rank $n - 1$ matrix.⁸ Both θ_c and Δ_c depend only on the variance-covariance matrix of returns, V_c , but not on the expected returns vector, $\tilde{\mu}_c$. The portfolio θ_c will play an important role in our analysis. We will refer to it as the *conditional global minimum variance portfolio*. Ideally, the investor would like to induce the manager to choose \tilde{x}_c^* . Note that, if the investor remains uninformed, then his optimal portfolio is $x_u^* = \theta_u + \tau \Delta_u \mu_u$, where θ_u is the global *unconditional* minimum variance portfolio and Δ_u is the appropriate matrix, and both are calculated using the unconditional variance-covariance matrix V_u .

The manager will obviously choose a portfolio that is optimal given her preferences and her compensation function. It is straightforward to show that the optimal portfolio for the manager given the compensation scheme in (2) is⁹

$$x = \frac{1}{h_1 + h_2} (h_2 b + h_1 \theta_c + \rho \Delta_c \tilde{\mu}_c). \quad (4)$$

8. These are given by $\theta_c = (e' V_c^{-1} e)^{-1} V_c^{-1} e$ and $\Delta_c = V_c^{-1} - (e' V_c^{-1} e)^{-1} V_c^{-1} e e' V_c^{-1}$, where e is the vector of ones.

9. Roll (1992) derives the locus of optimal portfolios for a manager whose compensation function is a special case of (2) with $h_1 = 0$.

If there exists a riskless security, whose return is normalized to zero, and if the manager is allowed to include it in her portfolio, then all the above holds with $\theta_c = 0$ and $\Delta_c = V_c^{-1}$.

The following result is obtained by comparing the manager's choice x with the optimal portfolio for the investor \tilde{x}_c^* .

PROPOSITION 1. The manager's portfolio choice is optimal for the investor if and only if (i) $h_1 + h_2 = \rho/\tau$ and (ii) either $h_2 = 0$ or $b = \theta_c$.

Thus, if a benchmark is used in the compensation of the manager and it is not set to the *conditional* global minimum variance portfolio, the resulting portfolio is not optimal for the investor. This includes the realistic case in which the benchmark is chosen to be the optimal portfolio given the unconditional distribution; that is, even if it is equal to x_u^* . Note that the result holds in the case where a riskless security exists if we interpret the global minimum variance portfolio to be the vector of zeros, meaning that the benchmark cannot involve any risky securities or the resulting portfolio is not optimal for the investor.

The intuition behind the proposition is as follows. For the manager's choice to be optimal for the investor it must be on the *conditional* mean-variance efficient frontier and must also be consistent with investor's preferences. It is well known that the mean-variance efficient frontier is spanned by any two efficient portfolios. One of these spanning portfolios can be taken as equal to the conditional global minimum variance portfolio θ_c . Note that this is the only portfolio on the conditional frontier whose composition the investor knows in advance, that is, which does not depend on private information observed only by the manager. The investor must guarantee that the first two terms in the manager's portfolio given by (4) are equal to the global minimum variance portfolio. This is brought about through condition ii of the proposition. The investor must also ensure that the last term is scaled appropriately to provide the optimal response for the investor. This is accomplished by condition i of the proposition. Note that, although our model assumes that the investor knows the characteristics of the manager's private information, the investor can obtain optimality without knowledge of V_c if he sets $h_2 = 0$; that is, he does not use benchmark-adjusted compensation.

To illustrate the economic magnitude of the loss associated with the use of an inappropriate benchmark, we now present an example. In this example we use the value weighted Standard and Poor's (S&P) 100 portfolio (the top 100 firms in the S&P 500 by market capitalization) as the benchmark. The variance-covariance matrix is estimated by factor-analytic techniques using 4 months of daily data from July to November 1996. For the purposes of the example we set expected returns so that the S&P 100 index is unconditionally mean-variance efficient with an expected excess return of 6%. We assume that both the investor and the manager have risk tolerance equal to 0.385. Given the assumed 6%

TABLE 1 Example of the Potential Costs Incurred in Using a Market Value Weighted Benchmark

Precision of Signals (1)	Value of Information with Optimal Response (%) (2)	Value of Information with Market Benchmark (%) (3)	Value of Information with Low Risk Benchmark (%) (4)	Loss with Market Benchmark (%) (5)	Loss with Low Risk Benchmark (%) (6)
.10	6.78	3.13	6.55	53.90	3.39
.09	6.11	2.45	5.88	59.85	3.76
.08	5.43	1.78	5.20	67.29	4.24
.07	4.74	1.09	4.51	77.10	4.85
.06	4.07	.41	3.84	89.93	5.65
.05	3.26	-.41	3.03	112.47	7.06
.04	2.58	-1.08	2.35	142.09	8.91
.03	1.90	-1.76	1.67	192.95	12.11
.02	1.22	-2.44	.99	300.78	18.85

expected excess return, this is the level of risk tolerance that leads one to hold the S&P 100 index if the S&P 100 were the available universe. The private information of the manager is assumed to be a noisy version of the firm-specific components in the return-generating process, where the variance of the noise terms is equal across stocks. The inverse of this variance measures the precision of the information.

In column 1 of table 1 we postulate various precision parameters, where the top entry represents the most precise signals and the precision is declining for lower entries. For each precision we calculate and report in column 2 the ex ante value of information to the investor (measured in certainty-equivalent percentage returns) attainable if he were using the information optimally to form his portfolio, when the alternative is to invest passively in the S&P 100 (the unconditional mean-variance efficient portfolio). For example, if the precision is 0.10, we find that the value of information with optimal response is 6.78%. This means that the information, if used optimally, allows the investor to increase his risk-adjusted expected return by 6.78%.

Now suppose that in this example the investor lets a manager with private information invest his wealth. Suppose further that the compensation of the manager is based only on the difference between the returns on the manager's portfolio and those of the value weighted S&P 100. According to proposition 1, since that benchmark is not equal to the conditional global minimum variance portfolio, it would not lead to the optimal portfolio choice. In column 3 of table 1 we report the ex ante value of the manager's services under these conditions. In the first row we see that, when the market benchmark is used, the information only increases the investor's risk-adjusted expected return by 3.13%. As reported in column 5, this is 53.9% less than the increase produced by the optimal response. The most striking finding occurs

when the precision of information is low. We see that when the precision is sufficiently low, the investor may lose more than 100% of the value of information when a market benchmark is used. This means that the risk-adjusted expected return the investor attains is less than what is attained by investing in the market portfolio. In such a circumstance the investor is better off not using the manager at all than using the manager with benchmark-adjusted compensation.

Column 4 of table 1 considers the case in which the benchmark is closer to the global minimum variance portfolio in that it has lower riskiness than that of the market portfolio. Specifically, we set the benchmark equal to the portfolio that would be chosen by an uninformed investor with risk tolerance equal to $\rho/4$. We see that the loss is significantly smaller in this case even for imprecise information. For the same parameters for which the loss with a market benchmark is 300.78%, the loss with the low risk benchmark is only 18.85%. Obviously, if the conditional global minimum variance portfolio is chosen as a benchmark, or no benchmark at all is used ($h_2 = 0$), then the full value of information would be realized and no loss incurred.

IV. Coordination Using a Benchmark Portfolio

So far we have assumed that the investor delegates all of his wealth to the portfolio manager. The use of benchmark portfolio is often discussed, however, in cases where the investor is attempting to coordinate the portfolio chosen by the manager with other holdings that he has. For example, the investor may hold some of his wealth in an index fund such as the S&P 500 and the rest may be given to one or more active portfolio managers to invest on his behalf. Often the benchmark in such a case is set equal to the passive portfolio held by the investor, with the motivation being that the benchmark signifies what the manager is attempting to “beat.”

To analyze this case, suppose that the investor gives to the manager a fraction π of his wealth and the rest is held in a passive portfolio whose risky security component is denoted by x_p . (If there is a riskless security, the elements of x_p do not need to sum to one.) *We do not make any assumption concerning the composition of x_p .*

The manager’s compensation is now given by

$$(h_1 x' \tilde{r} + h_2 (x - b)' \tilde{r}) \pi. \quad (5)$$

It is easy to show that the manager’s portfolio choice given this compensation is

$$x = \frac{1}{h_1 + h_2} (h_2 b + h_1 \theta_c + \frac{\rho}{\pi} \Delta_c \tilde{\mu}_c). \quad (6)$$

The following result generalizes proposition 1 to the case where π may be smaller than one. It is obtained simply by equating the optimal portfolio for the investor, \tilde{x}_c^* , with the investor's overall portfolio $\pi x + (1 - \pi)x_p$.

PROPOSITION 2. The investor obtains the optimal overall portfolio if and only if the following conditions hold:

$$h_1 + h_2 = \frac{\rho}{\tau}; \quad (7)$$

and, if $h_2 \neq 0$,

$$b = \theta_c + (\theta_c - x_p) \left(\frac{1 - \pi}{\pi} \right) \left(\frac{h_1 + h_2}{h_2} \right). \quad (8)$$

If $h_2 = 0$, then the investor obtains overall optimality if and only if $h_1 = \rho/\tau$ and either $\pi = 1$ or $x_p = \theta_c$.

Proposition 2 implies that it is possible for the investor to obtain overall optimality but that to accomplish this the chosen benchmark b must bear a particular relation to the portfolio that the investor holds on his own account x_p . Specifically, the difference between b and the conditional global minimum variance portfolio θ_c must be proportional to the difference between θ_c and x_p . Are there conditions under which it is optimal for the benchmark to equal the passive portfolio, that is, for b to equal x_p ? Proposition 2 shows that this only occurs when $x_p = \theta_c$. If, as is more typical, $x_p = x_u^*$ (the passive portfolio is equal to the portfolio that is optimal using the unconditional distribution), then the optimal benchmark does not equal the passive portfolio. Similarly, if there is a riskless security but the investor's passive portfolio includes some risky securities, then the optimal benchmark does not equal the passive portfolio since the optimal benchmark is the global minimum variance portfolio, which is the riskless security.

In Figure 1 we show how dramatic the difference can be between the benchmark which solves the coordination problem and the type of benchmarks that are typically used. The figure is based on the S&P 100 example used in the last section. Assume that $\pi = 0.5$. This means that half the investor's wealth is given to the active manager and the remainder is held in a passive portfolio. Assume that the passive portfolio is the S&P 100, which in our example is the unconditionally mean-variance efficient portfolio. Finally, assume that $h_1 = 0$. From proposition 2 we see that the appropriate benchmark under these circumstances is $2\theta_c - x_p$, or the difference between twice the minimum variance portfolio and the S&P 100. Figure 1 shows that this portfolio is quite different from the S&P 100, which is what would typically be used as a benchmark. (The optimal benchmark weights are given by

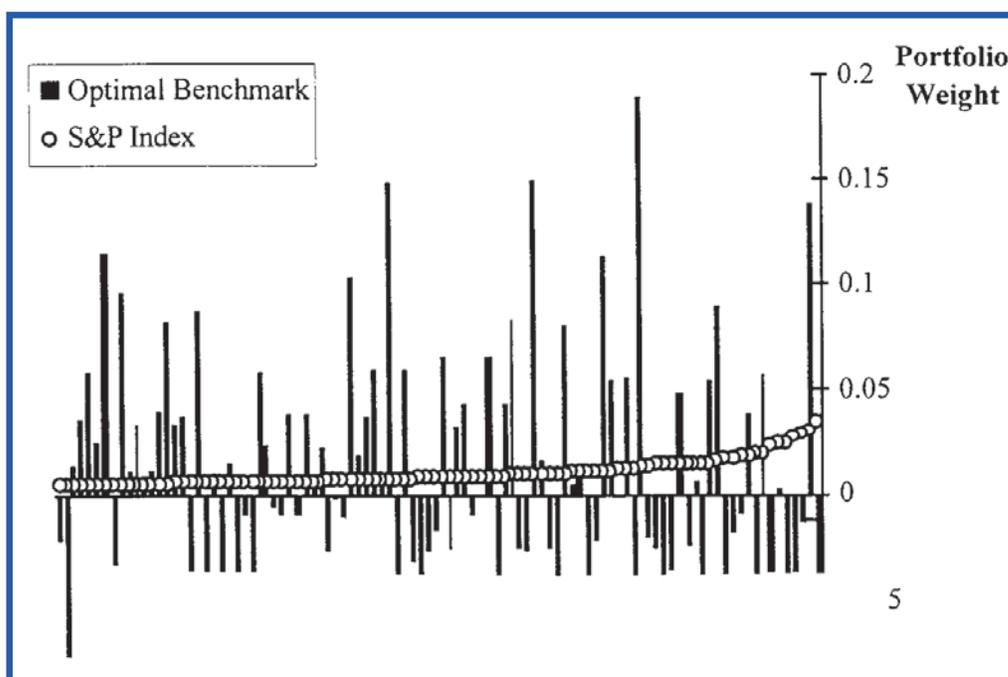


FIG. 1.—The 100 stocks in the S&P 100 (in ascending order by market capitalization).

the height of the bars, while the S&P 100 weights are given by the circles.)

We conclude that the use of standard benchmarks such as the market portfolio cannot be justified in the context of attempts to coordinate the manager's portfolio choice and align it with the investor's preferences. In the next few sections we take up issues related to risk sharing and examine the potential role of benchmarks in resolving various agency problems. For simplicity we will assume for the rest of the article that $\pi = 1$, that is, the manager controls the investor's entire portfolio. Most of the results below carry over in a straightforward manner to cases where $\pi < 1$.

V. Risk Sharing and the Cost of Compensating the Manager

The last two sections showed that benchmark-adjusted compensation schemes typically lead to suboptimal portfolio choices for the investor. However, this analysis ignored the cost of compensating the manager. It seems possible that, by shielding the manager from bearing the risk of the benchmark return, benchmark-adjusted compensation might lower the total cost of compensating the manager. This gain, if it exists, could in principle offset the losses identified in the last section. To examine this issue note that, in the absence of any contracting problems, the most cost-effective way to compensate the manager does not

involve benchmark-adjusted compensation. This is because, given our preference assumptions, optimal risk sharing requires that the investor and the manager share the risk of the total return proportionally, that is, they each receive a fixed fraction of the total return earned on the fund (see Wilson 1968). The manager's share of the risk depends on the level of her risk tolerance relative to that of the investor. Benchmark-adjusted compensation is obviously not consistent with the linear sharing rule required for optimal risk sharing.

While the linearity of the sharing rule depends on the assumption that both the investor and the manager have constant absolute risk aversion, we now argue that benchmark-adjusted compensation is not consistent with optimal risk sharing under *any* set of preferences represented by Von Neumann–Morgenstern state-independent utility functions. To see this let $U_M(\cdot)$ and $U_I(\cdot)$ be the utility functions for the manager and the investor, respectively. Then optimal risk sharing requires that, for all possible total returns earned on the fund R , the manager receive some fraction of the return α_R and the investor receive the rest, where α_R solves $U'_M(\alpha_R R) = \lambda U'_I[(1 - \alpha_R)R]$ for some positive constant λ . Benchmark-adjusted compensation does not achieve this, because the split between the manager and the investor should optimally depend only on the total return of the fund, but under benchmark-adjusted compensation the same total return of the fund can lead to different allocations to the manager and the investor.

VI. Benchmarks and the Effort Incentive Problem

The results of the last section show that there are potentially large costs involved in using benchmarks such as the market portfolio in compensation contracts. In this and the next section we ask whether these costs might be offset by some benefits that benchmarks may offer in solving various agency problems, as often alluded to in discussions of benchmarks.¹⁰

Suppose that the manager must expend costly effort to generate private information. After the effort choice is made and effort is expended, the manager observes the signal and makes the investment decision. Finally, after returns are realized, the manager obtains the compensation from the investor. The issue we address is whether the use of certain benchmarks in the compensation structure can enhance the manager's incentives to expend the appropriate amount of effort from the investor's perspective.

We continue to assume that the manager can obtain information about the idiosyncratic terms in the factor model of returns. However,

10. See, e.g., Rennie and Cowhey (1990).

we now assume that to obtain this information the manager must expend costly effort. Intuition from the moral hazard literature may suggest that the appropriate incentive contract should expose the manager to risks that are related to her effort choice but that she should not be exposed to risks that are outside of her control, unless this is important for risk sharing. This seems to imply in this example that the manager's compensation should depend on the idiosyncratic returns of the securities but not on the factor realizations. It also seems that the way to accomplish this is to make the manager's compensation depend on the performance of her portfolio relative to an appropriately chosen benchmark. It would appear that by netting out the benchmark most of the factor risk can be removed, leaving the manager essentially exposed only to idiosyncratic risk. It turns out, however, that, given the way the manager responds to the benchmark adjustment, this does not happen.

The following result shows that in the case where there is a riskless security, the existence and composition of the benchmark portfolio are completely irrelevant to the effort-incentive problem. Whatever incentives can be created with a particular benchmark, the same incentives can be created with a different benchmark or without the benchmark at all. (As we will discuss later, the same is essentially true without a riskless security.) In this result we do not assume that the riskless return is zero since the manager's compensation, and therefore the effort choice, may depend on the riskless rate. Denote the riskless return by r_f . Making the appropriate adjustments, suppose the manager's compensation is given by

$$h_1[x'\tilde{r} + (1 - e'x)r_f] + h_2[(x - b)'\tilde{r} + e'(x - b)r_f], \quad (9)$$

where x is the manager's chosen portfolio of risky securities, b is a benchmark portfolio, and e is the vector of ones. (The components of both x and b are not restricted to sum to one.) We have the following result:

PROPOSITION 3. Under the above conditions, the manager's effort choice is independent of the composition of b and of the coefficients h_1 and h_2 .

Proof. We know that, when there is a riskless security, the optimal portfolio choice for the manager under the compensation scheme given above is

$$x = \frac{1}{h_1 + h_2} (h_2 b + \rho V_c^{-1} \tilde{\mu}_c). \quad (10)$$

Using this, it follows that

$$h_1 x = \frac{h_1 h_2}{h_1 + h_2} b + \frac{h_1 \rho}{h_1 + h_2} V_c^{-1} \tilde{\mu}_c, \quad (11)$$

and

$$h_2(x - b) = -\frac{h_1 h_2}{h_1 + h_2} b + \frac{h_2 \rho}{h_1 + h_2} V_1^{-1} \tilde{\mu}_c. \quad (12)$$

The manager's total exposure to the risky security returns is therefore

$$[h_1 x + h_2(x - b)]' \tilde{r} = \rho \tilde{\mu}_c' V_c^{-1} \tilde{r}. \quad (13)$$

Note that this does not depend on h_1 , h_2 , or b . The manager's effort affects V_c and the signals received that lead to the forecast $\tilde{\mu}_c$, but since the effect of these on the manager's compensation is independent of h_1 , h_2 , and b , it follows that the effort choice is independent of these as well. Q.E.D.

It should be noted that the irrelevancy of the *composition* of the benchmark is quite a general result that does not depend on our parametric specification. If the manager's compensation is *any function* of $[h_1 x + h_2(x - b)]' \tilde{r}$ for some coefficients h_1 and h_2 , then the manager can achieve exactly the same payoff realization no matter what the composition of b is. Whatever the optimal portfolio for the manager is for one given benchmark, there is another portfolio that can produce exactly the same payoff for the manager if the benchmark composition is changed.

To understand the intuition behind proposition 3, consider a general agency problem in which a principal hires an agent to produce output in a project, and assume that the total output of the project is $\ell(\tilde{\omega}_a + \tilde{\delta})$ where $\tilde{\omega}_a$ and $\tilde{\delta}$ are random variables and ℓ is a positive constant. Suppose the distribution of $\tilde{\omega}_a$ depends on the agent's unobservable and costly effort level a , while the distribution of $\tilde{\delta}$ is not influenced by anything the agent does. Suppose that the principal pays the agent a fraction k of the total output, that is, $k\ell(\tilde{\omega}_a + \tilde{\delta})$. Now consider two different cases. In one, the constant ℓ is outside the manager's control. In this case, by making k larger the principal can make the agent more sensitive to the effects of increased effort. Thus, the choice of k potentially affects the effort choice of the manager. In the second case the agent has complete and costless control over ℓ . Now it is clear that the principal has no ability to affect the agent's choice of effort through choice of k because any attempt by the principal to sharpen incentives by doubling k would just induce the agent to halve the level of ℓ . Clearly, our setting is analogous to the second case since the manager in our model has complete control over the scale of her response to her signals.

Note that, in the first case above, with ℓ outside the manager's control, if k is increased, incentives are sharpened, but the agent also bears more of the risk due to the noise term $\tilde{\delta}$, which by assumption is outside the agent's control. If the principal can observe $\tilde{\delta}$ (or a variable corre-

lated with it), then the cost of providing the sharper incentives produced by increasing k can be reduced by subtracting off some part of the risk created by $\tilde{\delta}$. However, if the agent has complete control over ℓ , then it is not possible to sharpen incentives by increasing k , and deviating from optimal risk sharing is not useful for the purpose of providing incentives. This is indeed the situation in our portfolio management context. The return on an appropriately specified benchmark portfolio can be interpreted as noise in the same sense as $\tilde{\delta}$, but given the portfolio manager's ability to control the scale of her response, the investor has no ability to provide incentives, and therefore deviating from optimal risk sharing is not appropriate. Benchmark-adjusted compensation therefore has no constructive role to play in this context.

Proposition 3 and the discussion above are predicated on the existence of a riskless security. When no riskless security exists, the irrelevancy of the composition of the benchmark remains true, but it is possible for the effort choice to depend on the coefficients h_1 and h_2 . To see this, note that the manager's total exposure to the risky securities return becomes

$$h_1 x' \tilde{r} + h_2 (x - b)' \tilde{r} = h_1 \theta_c' \tilde{r} + \rho \tilde{\mu}'_c \Delta_c \tilde{r}. \quad (14)$$

Both the conditional global minimum variance portfolio θ_c and the matrix Δ_c depend on the conditional variance-covariance matrix V_c , which is affected by the effort choice of the manager. As h_1 varies, the importance of the return on the global minimum variance portfolio changes, and this might have some effect on incentives. Numerical examples show, however, that this effect is very minor and that having $h_2 > 0$ actually *weakens* the incentive of the manager to expend effort. In all cases, if a benchmark is used, the composition of the benchmark is completely irrelevant.

VII. On the Use of Benchmarks to Infer the Manager's Skill

Up to this point we have assumed that the investor knows how skillful the manager is and that he can correctly evaluate the manager's information.¹¹ In practice investors often have significant uncertainty about the skill of portfolio managers. Indeed, a great deal of effort is spent in attempts to assess and measure portfolio managers' performance given a time series of historical returns on their and other portfolios.

11. Recall, however, that, even if the investor does not know precisely the precision of the manager's information signals, optimality can still be achieved by simply setting $h_1 = \rho/\tau$ and $h_2 = 0$ in the compensation function, i.e., without using a benchmark. Of course the investor is only willing to hire the manager under this contract if he is convinced that the manager has some skill and is not responding to spurious information signals. For a model that shows how a contract can be used by an investor to screen out inferior managers, see Bhattacharya and Pfleiderer (1985).

In this section we consider three aspects of the problem. First (in Subsection *A*), we look at the pure inference problem and ask whether benchmark-adjusted returns are a sufficient statistic for assessing the value of the manager's information given the returns on individual securities and the manager's returns. We show that they are not, even when the precision of the manager's information depends on only a single parameter. Second (in Subsection *B*), we examine whether benchmark-adjusted compensation enhances the investor's ability to assess the manager's skills and screen out bad managers and how the choice of benchmark affects this inference. We show that, similar to the case of effort incentives, benchmarks and especially their compositions are largely irrelevant in this regard. Finally (in Subsection *C*), we consider the possibility that explicit benchmark adjustments are made to counteract the potential response of the manager to implicit benchmarks used, for example, in ex post assessments of her performance. We find that explicit benchmarks can be chosen to do so but that the benchmarks that have the desired effects are far different than the ones used in practice.

A. The Bayesian Inference Problem

If the distributional parameters of the manager's information, which obviously are related to the manager's skill, are not known to the investor or to other potential investors, then ex post observations of returns can be used to attempt to infer these parameters. Suppose ex post observable variables include the returns on all individual securities and the returns on the manager's portfolio. The question we examine is whether benchmark-adjusted returns (i.e., the difference between the manager's return and those on a prespecified benchmark portfolio) can be a sufficient statistics in this inference problem.

The exact nature of the inference problem depends of course on the unknown parameters and their assumed joint distributions with the observables. In the appendix we develop a formal model in which there is one unknown precision parameter. We show that, even in this simple one-parameter setting, the appropriate Bayesian updating rule is not generally expressible in terms of the difference between the manager's return and the return on a benchmark.

To see the argument intuitively, let b be an arbitrary benchmark portfolio and consider a situation where the ex post return on the manager's portfolio is equal to the return on the benchmark, that is, $(x - b)'r = 0$. If $(x - b)'r$ is a sufficient statistic, then the inference one draws about the manager's skill when $(x - b)'r = 0$ must be the same no matter what the realization of returns r . But this cannot be true. If the realized returns on all securities turn out to be the same, then one has learned nothing about a manager's skill since the return on all portfolios is the same. In this case one would not update the priors at all. However, if

many of the realized security returns are far from their unconditional means, we would expect a well-informed manager to have received signals that would indicate this and have adjusted the portfolio accordingly. Under these conditions one clearly would adjust one's prior for a manager that posts the same return as the benchmark.

The above argument shows that assessing a manager's skill *solely* on the basis of benchmark-adjusted returns is not consistent with any Bayesian updating rule. It is possible that assessments based on performance relative to a benchmark are good "rules of thumb" given the complexity of the Bayesian inference problem. However, examining the extent to which this might be true is beyond the scope of this article.

B. An Irrelevancy Result

In this section we ask whether the particular parameters the investor can set in designing a benchmark-adjusted compensation scheme can affect his ability to infer the manager's skill from ex post returns. For example, is there a particularly appropriate benchmark choice (e.g., one that somehow captures the type of information the manager claims to have) that can enhance this assessment? The following result shows that, similar to the case of effort incentives, the ex post inference that one draws about the manager's ability is generally independent of the exact form of the manager's compensation. For this result we assume that there is a riskless security whose return is to zero.

PROPOSITION 4. Assume that a riskless security is available and that the investor has some prior assessment of the precision of the manager's information. Then the investor's posterior belief about the precision given the realized returns on the manager's portfolio and the return on individual securities is independent of the composition of b and of the coefficients h_1 and h_2 .

Proof. From our results above we know that

$$\tilde{x}'_1 \tilde{r} = \frac{h_2}{h_1 + h_2} b' \tilde{r} + \frac{\rho}{h_1 + h_2} \tilde{\mu}'_c V_c^{-1} \tilde{r}. \quad (15)$$

The distribution of the manager's signals affects $\tilde{x}'_1 \tilde{r}$ only through $\tilde{\mu}'_c V_c^{-1} \tilde{r}$, and this is easily recovered from the observed returns since

$$\frac{(h_1 + h_2)}{\rho} \tilde{x}'_1 \tilde{r} - \frac{h_2}{\rho} b' \tilde{r} = \tilde{\mu}'_c V_c^{-1} \tilde{r}. \quad (16)$$

Thus for any h_1 , h_2 , and b (with $h_1 + h_2 \neq 0$), the posterior distribution is based on the same statistic $\tilde{\mu}'_c V_c^{-1} \tilde{r}$, and so it is independent of the compensation. Q.E.D.

Propositions 3 and 4 are obviously related. The key to understanding them is that the manager is able to change his portfolio in response to the compensation structure. This response "undoes" any possible ef-

fect of the particular benchmark and coefficients in the compensation structure on the inference concerning the manager's skill.¹² Efforts by the investor to choose an appropriate benchmark so as to be able to learn better the type of information possessed by the manager are therefore fruitless.

C. *A Model with Implicit and Explicit Benchmarks*

The analysis above assumes that, even if the manager's skill is unknown, the compensation structure is still given in the general form of (2). To the extent that subsequent inference is made concerning the manager's skill, this has so far not entered the compensation structure, and so it has not affected her portfolio choices. In reality, however, inferences made about the manager's skill typically do affect the manager's compensation since the amount of money the manager is given to invest in the future is likely to depend on the assessment of her skill. We now propose a simple, reduced-form model in which the manager's compensation is potentially based on two benchmarks. In addition to the "explicit" benchmark set by the investor, there is another "implicit" benchmark adjustment, this one motivated by the ex post inference regarding the manager's skill. The manager responds to both adjustments in choosing her portfolio.

We denote by b_e the explicit benchmark set by the investor in the manager's compensation and by b_i the implicit benchmark used in the inference. We postulate that the manager receives proportionally more funds to manage in the future as a function of her performance relative to the implicit benchmark. For analytical simplicity we assume that the relation between the manager's compensation and performance relative to the implicit benchmark is linear. The total compensation of the manager is therefore

$$h_1 \tilde{x}' \tilde{r} + h_{2e} (\tilde{x} - b_e)' \tilde{r} + h_{2i} (\tilde{x} - b_i)' \tilde{r}, \quad (17)$$

where \tilde{x} is the manager's portfolio choice. If it is assumed that a riskless security whose return is zero is available, then it is straightforward to show that the manager's response to the compensation structure given by (17) is

$$\tilde{x} = \frac{1}{h_1 + h_{2e} + h_{2i}} (h_{2e} b_e + h_{2i} b_i + \rho V_c^{-1} \tilde{\mu}_c). \quad (18)$$

12. Note that, if a riskless security does not exist, there is uncertainty about the conditional global minimum variance portfolio θ_c since this depends on the manager's ability. The composition of the benchmark continues to be irrelevant in this case, but now the inference does depend on the relative magnitude of h_1 and h_2 . However, it can be shown that the effect of this is likely to be quite small.

We now assume that, taking b_i and h_{2i} as given, the investor sets h_1 , h_{2e} , and b_e to attempt to induce a portfolio choice by the manager that is optimal for the investor. An argument similar to that used for proving proposition 2 shows that the investor can achieve the optimal portfolio only if $h_{2e} \neq 0$, that is, if he uses an explicit benchmark. Optimality is achieved by setting

$$h_1 + h_{2e} = \frac{\rho}{\tau} - h_{2i} \quad (19)$$

and

$$b_e = -\frac{h_{2i}}{h_{2e}} b_i. \quad (20)$$

Thus, if the investor is aware that the manager's portfolio choice responds to the presence of the implicit contract b_i , he can design the explicit benchmark b_e , as well as the coefficients h_1 and h_{2e} in the compensation function, so as to "undo" the effect of the implicit benchmark on the manager's portfolio choice. Given b_i , the explicit benchmark b_e is a necessary part of the compensation contract that induces the optimal portfolio choice for the investor.¹³ This is similar to the role of the benchmark in coordinating the manager's portfolio with the investor's other passive holdings, as analyzed in Section IV above.

We have motivated the presence of the implicit benchmark by postulating that investors use performance relative to a benchmark in trying to assess the manager's skill. As we noted earlier in Subsection VIIA, this is at best an approximation to the true solution of the Bayesian inference problem faced by investors. As such, it is difficult to determine what a "reasonable" implicit benchmark should be. In practice, the benchmark used in inference about managers' skill is often the market portfolio or the optimal unconditional portfolio, with the motivation that this is the portfolio the manager must "beat" if she is informed. Note that the optimal unconditional portfolio $\rho V_u^{-1} \mu_u$ is the unique portfolio that has the following property: if b_i is equal to that portfolio and the investor sets b_e , h_{1e} , and h_{2e} optimally, then an uninformed manager will have the same returns as b_i . In other words, only if the implicit benchmark b_i is set equal to $\rho V_u^{-1} \mu_u$ will an uninformed manager perform no better and no worse than b_i .

Suppose in fact that b_i is set equal to $\rho V_u^{-1} \mu_u$. Then (20) becomes

$$b_e = \frac{-h_{2i}}{h_{2e}} \rho V_u^{-1} \mu_u. \quad (21)$$

13. Note that the same analysis would apply if b_i was another "explicit" benchmark (set, e.g., by another investor) in the manager's total compensation.

Thus, in order to induce optimal portfolio choice for the investor, in this case b_e must be proportional to the *negative* of the market portfolio. This does not conform to what is typically observed in compensating portfolio managers. In particular, if b_e is also set equal to the market portfolio, it can be shown that a substantial loss of value occurs in this model.

VIII. Benchmarks When the Manager's Risk Tolerance Is Unknown

So far we have assumed that the manager's preferences are known to the investor. Recall from propositions 1 and 2 that in order to align the manager's portfolio choice and particularly its risk attributes with the desired portfolio for the investor, the coefficients h_1 and h_2 must be set so that their sum is equal to the ratio of the manager's risk tolerance and that of the investor. If the manager has a higher risk tolerance than anticipated by the investor (and used in the compensation function), the compensation contract will induce the manager to take on too much risk, and, conversely, if the investor overestimates the manager's risk tolerance, the manager's choice will be overly conservative.

In his discussion of the effect of benchmark portfolios Roll (1992) examines the possibility that the investor imposes constraints on the beta of the manager's portfolio. Intuitively, constraints on some quantifiable risk attributes of the manager's portfolio may well control the risk characteristics of the portfolio chosen by the manager. Numerical calculations in our example of Section III show that beta constraints of this sort can substantially reduce (but not eliminate) the loss created by the benchmark.¹⁴ Note that the fact that beta constraints may be useful does not explain why benchmark-adjusted compensation should be used instead of, for example, compensation that depends only on the total returns on the manager's portfolio. Indeed, it is the presence of the benchmark that causes the risk return trade-offs faced by the manager to be distorted relative to those of the investor. The question arises whether the use of benchmark can be justified in a context in which it is appropriate to constrain the riskiness of the manager's choice for some other reason.

The following proposition shows that in fact there is no reason to use a benchmark in addition to a set of risk constraints when the manager's risk tolerance is unknown; the benchmark does not add anything and in fact can make the investor worse off. The proposition compares two compensation schemes. In both the manager's portfolio must satisfy a set of risk constraints. In one scheme the manager's compensa-

14. A potential problem with the beta constraint is that the beta of the optimal conditional portfolio is not fixed but rather varies with the information realization. However, with a large number of securities this is not a substantial problem.

tion is based only on the total return on her portfolio. In the second scheme compensation is benchmark-adjusted, and the benchmark satisfies the same set of risk constraints. The proposition shows that, except in one knife edge case (in which the benchmark is the constrained conditional global minimum variance portfolio and in which we show that the investor is indifferent between these schemes), the investor is strictly better off if the first scheme is used, that is, no benchmark adjustment is made.

PROPOSITION 5. Let C be an $n \times j$ matrix and let q be a $j \times 1$ vector. Assume that the manager's risk tolerance ρ is unknown but is drawn from a symmetric distribution with mean $\bar{\rho}$. Let $h = \bar{\rho}/\tau$ and let U_1 and U_2 be the expected ex ante utility level achieved by the investor under the following two compensation schemes respectively.

- (1) The manager receives $hx'\tilde{r}$ but is required to choose x such that $C'x = q$.
- (2) The manager receives $h(x - b)'\tilde{r}$ where b is a benchmark that satisfies $C'b = q$. The manager is required to choose x such that $C'x = q$.

Then $U_1 \geq U_2$ for all benchmarks satisfying $C'b = q$. Moreover, $U_1 = U_2$ only if $b = b^*$, where b^* is the solution to $\min_y y'V_c y$ subject to $C'y = q$.

Proof. See the appendix.

This is yet another negative result concerning the use of benchmarks.

IX. Concluding Remarks

In this article we have presented a number of results concerning the practice of using benchmark adjustments in the compensation of portfolio managers. We have examined many obvious and straightforward possible explanations for the use of benchmarks in compensation, including risk sharing, coordination, effort incentives, and screening. Yet we have been unable to find (at least within our model) a setting in which such contracts emerge as the solution to a contracting problem. It is possible that benchmark adjustments arise naturally in other models of portfolio management. In our model, however, our results lead us to question the value of using such an adjustment.

While we have focused on the case of one investor interacting with one manager, our approach and some of our results have immediate implications for the more realistic case in which many investors interact with many managers. For example, it is possible to extend proposition 2 to the case in which the manager invests on behalf of multiple investors, each of whom may set a different benchmark. Not surprisingly, obtaining overall optimality for all investors imposes extremely strong

conditions on the passive portfolios of the investors, on the coefficients in the compensation function and on the various benchmarks used. For the case of multiple managers working with the same investor it is not possible to obtain a positive result analogous to proposition 2. That is, as long as the investor allocates a positive fraction of his wealth to more than one informed manager and managers' information signals are not perfectly correlated, it is not possible in general to coordinate the managers' portfolios in such a way that the overall portfolio of the investor is optimal. This is true with and without the use of benchmark portfolios. To see this, note that the overall optimal portfolio depends on the conditional variance and the expected return vector *given all the private information of all the managers*. However, under any compensation structure of the type given in (2), each manager calculates her response using the conditional distribution based only on her information. The compensation structure we assume is not capable of setting the appropriate risk-return trade-offs for achieving overall optimality.

Note that a simple solution of the coordination problem in the case of one investor interacting with one or many managers is to tie each manager's compensation to the *overall* portfolio return of the investor (rather than tying each manager's compensation only to the return on that manager's portfolio). This would immediately align the managers' preferences with those of the investor. A problem with this is that it significantly weakens the managers' incentives to collect information and possibly creates other agency problems. Thus, interesting trade-offs may arise in this setting.

In our analysis we have restricted the compensation structure to depend only on the returns of the manager's portfolio but not on the manager's actual portfolio choices. If portfolio choices are observable and if it is possible to construct explicit or implicit compensation contracts based on them, then the investor should be able to obtain better results in attempting to induce the appropriate portfolio choice. In the context of the screening problem and performance measurement, it is obviously easier to infer the manager's information structure based on the observation of her portfolio choices. In general, better coordination across managers is likely to be achievable using such observations.

The benchmark portfolios considered in our model are "passive," that is, they have fixed weights that do not depend on any private information. This property is often advocated as desirable for benchmark portfolios (see, e.g., Bailey and Tierney 1993). In the spirit of the literature on tournaments, however, one might wish to consider a benchmark whose return is the average return earned by a set of potentially informed, "active" managers.

Another assumption we have made is that there are no frictions such as short sales constraints (or costs) that affect the portfolio choices of

the manager or the investor. Such restrictions may affect our results in the sense that, for example, different benchmarks would cause the portfolio choices of a manager to hit the constraints with differing frequency.

The above comments suggest a number of directions for future research. We suspect, however, that in order to justify the use of benchmark portfolios in compensation one would have to develop a model substantially different from ours. Absent such a development, it seems that the use of benchmarks in compensation remains problematic.

Appendix

Demonstration That a Benchmark-Adjusted Return Is Not a Sufficient Statistic

Let us fix a linear compensation contract given by (2) and assume that a manager chooses portfolios in response to this contract. We assume that the manager observes an n -dimensional signal \tilde{s} equal to $\tilde{\epsilon} + \tilde{\xi}$ where $\tilde{\xi}$ has precision matrix (inverse of the variance-covariance matrix) equal σQ . To make the inference problem as simple as possible, we assume that Q is a positive-definite matrix that is known by the investor, while σ is a nonnegative scalar that the investor does not know. In other words, the manager's skill is measured by the one-dimensional parameter σ , where a higher σ means more precise information. We include the possibility that the manager is uninformed, in which case the σ is zero. We assume that the investor has a prior distribution over the value of σ given by $g_0(\sigma)$. We will treat the distribution as a continuous one (in which case $g_0(\sigma)$ is a density), but the argument carries over to cases where the distribution has mass points at various values of σ . Finally we assume that there is a riskless security with return r_f .

In analyzing the inference problem we assume that the investor observes the realized return on the manager's fund $x'r + (1 - x'e)r_f$, which is the realization of $\tilde{x}'\tilde{r} + (1 - x'e)r_f$ and the realized return of each of the securities in the market (i.e., the vector r of realizations of \tilde{r}). Let $\phi(z|\sigma, r)$ be the density of the manager's return distribution given that the quality of the manager's information is σ that the realized security returns are r . To calculate this density we must consider each vector of realizations x that satisfies $x'r + (1 - x'e)r_f = z$. For each of these x there is a unique μ_c (vector of conditional expectations) that gives rise to it, and for each μ_c there is a unique signal vector that supports it. This means that for a given σ there is a functional relation between x and s . We let $s = S(x, \sigma)$ be the signal that leads to x when the manager has ability σ . Thus, we find the density of the manager's observed signal vector given ex post security returns and a quality parameter σ . Let this be $\psi(s|\sigma, r)$. The value of $\phi(z|\sigma, r)$ is obtained by integrating over all of the signal vector densities $\psi(s|\sigma, r)$ that are associated with the x 's that satisfy $x'r + (1 - x'e)r_f = z$. In other words,

$$\phi(z|\sigma, r) = \int_{\{x: x'r + (1-x'e)r_f = z\}} \psi(S(x, \sigma)|\sigma, r) dx. \quad (A1)$$

Finally, the posterior distribution is given by $g_1(\hat{\sigma}|x'r, r)$, where

$$g_1(\hat{\sigma}|x'r + (1 - x'e)r_f, r) = \frac{\phi(x'r + (1 - x'e)r_f|\hat{\sigma}, r)g_0(\hat{\sigma})}{\int \phi(x'r + (1 - x'e)r_f|\sigma, r)g_0(\sigma)d\sigma}. \quad (\text{A2})$$

We now argue that it is not possible for a benchmark-adjusted return to be a sufficient statistic. For this to be true, there must exist a benchmark b and a function $g_1^*(\cdot|\cdot)$ such that $g_1^*(\hat{\sigma}|(x - b)'r + r_f(b - x)'e) = g_1(\hat{\sigma}|x'r + (1 - x'e)r_f, r)$ for all $\hat{\sigma}$ and all r . To show that such a benchmark cannot be found, we will essentially show that there exists a set of outcomes over which the appropriate posterior distribution is not constant even though the benchmark-adjusted return is.

To begin we let b be any benchmark. First note that whenever all risky securities have the same return as the riskless security, that is, $r = r_f e$, the posterior distribution must be the same as the prior. This is because the returns on all portfolios are the same whenever $r = r_f e$, and so no information about the quality of the manager's signal is revealed. Thus $g_1(\sigma|r_f, r_f e) = g_0(\sigma)$ for all σ , and in particular $g_1(0|r_f, r_f e) = g_0(0)$. Note also that when $r = r_f e$, the benchmark-adjusted return on all portfolios is zero. Now let x_u be the optimal portfolio for an uninformed manager, and let δ be any nonzero n -vector such that δ is orthogonal to both x_u and b . If the vector of returns earned on the risky securities is $r_f e + \delta$, an uninformed manager will have a return of r_f and a benchmark-adjusted return of zero. However, when $r = r_f e + \delta$, an informed manager would have in expectation a different return than the uninformed manager since the informed manager's portfolio responds to information. Thus, if $r = r_f e + \delta$ and the manager has return of r_f , the posterior that the manager is uninformed is *higher* than the prior, that is, $g_1(0|r_f, r_f e) > g_0(0)$. However, if the manager is in actuality uninformed and $r = r_f e + \delta$, then her benchmark-adjusted return is zero, as is the case when $r = r_f e$. Any inference based only on the benchmark-adjusted return will not distinguish the case where $r = r_f e$ from the case where $r = r_f e + \delta$ even though $g_1(0|r_f, r_f e + \delta) \neq g_1(0|r_f, r_f e)$.

PROOF OF PROPOSITION 5. Under compensation scheme 1, the manager will set his portfolio choice equal to

$$x = Z_0 q + \frac{\rho\tau}{\rho} Z_1 \mu_c,$$

where $Z_0 = V_c^{-1} C(C' V_c^{-1} C)^{-1}$ and $Z_1 = V_c^{-1} - V_c^{-1} C(C' V_c^{-1} C)^{-1} C' V_c^{-1}$. Under compensation scheme 2 the manager will choose his portfolio such that

$$x = b + \frac{\rho\tau}{\rho} Z_1 \mu_c,$$

where Z_1 is as above. Since it can be shown that $b^* = Z_0 q$, it follows that, when $b = b^*$, the same portfolio is chosen under both schemes, and the investor attains the same utility under each. If $b \neq b^*$, then $b = b^* + \eta$ where η is orthogonal to all of the column vectors in C . Thus if x is the portfolio chosen under scheme (1), $x + \eta$ is the portfolio chosen under scheme (2). Now define $U_1(\mu_c, \bar{\rho} + \delta)$ to be the expected utility of the investor achieved under scheme 1 conditional on the manager's expectation of security returns being μ_c and conditional on the

manager's risk tolerance being $\rho = \bar{\rho} + \delta$. Define $U_2(\mu_c, \bar{\rho} + \delta)$ in a similar manner for compensation scheme 2. It is straightforward to show that

$$U_1(\mu_c, \bar{\rho} + \delta) = - \exp \left\{ - \frac{1}{\tau} \left[\mu_c' Z_0 q + \tau \left(\frac{1}{2} - \frac{\delta^2}{2\rho^2} \right) \mu_c' Z_1 \mu_c - \frac{1}{2\tau} q'(C' V_c^{-1} C)^{-1} q \right] \right\}$$

and

$$U_2(\mu_c, \bar{\rho} + \delta) = - \exp \left\{ - \frac{1}{\tau} \left[\mu_c' Z_0 q + \tau \left(\frac{1}{2} - \frac{\delta^2}{2\rho^2} \right) \mu_c' Z_1 \mu_c - \frac{1}{2\tau} q'(C' V_c^{-1} C)^{-1} q - \frac{\delta}{\rho} \eta' \mu_c - \frac{1}{2\tau} \eta' V_c \eta \right] \right\}.$$

Now fix t at some value and consider two realizations of δ , $\delta = t$ and $\delta = -t$. Since the distribution of ρ is symmetric, these two values of δ are equally likely. However,

$$\frac{1}{2} (U_2(\mu_c, \bar{\rho} + t) + U_2(\mu_c, \bar{\rho} - t)) < \frac{1}{2} (U_1(\mu_c, \bar{\rho} + t) + U_1(\mu_c, \bar{\rho} - t)).$$

This is because when we condition on $\tilde{\delta}^2$ being equal to a constant, $U_1(\mu_c, \bar{\rho} + \tilde{\delta})$ can be written as $f(a)$ and $U_2(\mu_c, \bar{\rho} + \tilde{\delta})$ as $f(a + \tilde{\omega})$, where $f(\cdot)$ is a concave function, a is a constant, and $\tilde{\omega}$ is a random variable with a negative mean equal to $-\eta' V_c \eta / 2\tau$. (Note that $-\eta' V_c \eta / 2\tau$ is strictly negative since V_c is positive definite and by assumption $\eta \neq 0$.) Since this is true for all values of t and all μ_c , it follows that

$$E(U_1(\tilde{\mu}_c, \tilde{\rho})) > E(U_2(\tilde{\mu}_c, \tilde{\rho})).$$

This concludes the proof. Q.E.D.

References

- Admati, A. R., and Ross, S. A. 1985. Measuring investment performance in a rational expectations equilibrium model. *Journal of Business* 58 (January): 1–26.
- Bailey, J. V. 1990. Some thoughts on performance-based fees. *Financial Analysts Journal* 4 (July–August): 31–40.
- Bailey, J. V., and Tierney, D. E. 1993. Gaming manager benchmarks. *Journal of Portfolio Management* 19 (Summer): 37–40.
- Bhattacharya, S., and Pfleiderer, P. 1985. Delegated portfolio management. *Journal of Economic Theory* 36 (June): 1–25.
- Brennan, M. J. 1993. “Agency and asset pricing,” pp. 1–33. Working Paper no. 6-93. Los Angeles: University of California, Los Angeles, Anderson Graduate School of Management.
- Gastineau, G. L. 1994. Beating the entity benchmarks. *Financial Analysts Journal* 50 (July–August): 6–11.
- Grinblatt, M., and Titman, S. 1989. Adverse risk incentives and the design of performance-based contracts. *Management Science* 35 (July): 807–22.
- Hansen, L. P., and Richard, S. F. 1987. The role of conditioning information in deducing

- testable restrictions implied by dynamic asset pricing models. *Econometrica* 55 (May): 587–613.
- Heinkel, R., and Stoughton, N. M. 1994. The dynamics of portfolio management contracts. *Review of Financial Studies* 7 (Summer): 351–87.
- Huddart, S. 1994. Reputation and performance fee effects on portfolio choice by investment advisers, pp. 1–38. Working Paper no. 1282. Stanford, Calif.: Stanford University, Graduate School of Business.
- Rennie, E. P., and Cowhey, T. J. 1990. The successful use of benchmark portfolios: A case study. *Financial Analysts Journal* 46 (September–October): 18–26.
- Roll, R. 1992. A mean/variance analysis of tracking error. *Journal of Portfolio Management* 8, no. 4 (Summer): 13–22.
- Sharpe, W. F. 1981. Decentralized investment management. *Journal of Finance* 36 (May): 217–34.
- Wilson, R. B. 1968. On the theory of syndicates. *Econometrica* 36 (January): 119–32.