

ISSN 1503-299X


WORKING PAPER SERIES

No. 2/2006

PORTFOLIO CHOICE WHEN MANAGERS CONTROL RETURNS

Egil Matsen

Department of Economics

 Norwegian University of Science and Technology

N-7491 Trondheim, Norway

www.svt.ntnu.no/iso/wp/wp.htm

Portfolio Choice when Managers Control Returns*

Egil Matsen **

Abstract

This paper investigates the allocation decision of an investor with two projects. Separate managers control the mean return from each project, and the investor may or may not observe the managers' actions. We show that the investor's risk-return trade-off may be radically different from a standard portfolio choice setting, even if managers' actions are observable and enforceable. In particular, feedback effects working through optimal contracts and effort levels imply that expected terminal wealth is nonlinear in initial wealth allocation. The optimal portfolio may involve very little diversification, despite projects that are highly symmetric in the underlying model. We also show that moral hazard in one of the projects need not imply lower allocation to that project. Expected returns are generally lower than under the first-best, but the optimal contract shifts more of the idiosyncratic risk in the hidden action project to the manager in charge of it. The minimum-variance position of the investor's (net) terminal wealth would in most cases involve a portfolio shift towards the hidden action project, and there are plausible cases where this would dominate the overall effect on the second-best optimal portfolio when comparing with the first-best.

Keywords: Portfolio choice, diversification, optimal contracts

JEL Classifications: D81; D82; G11

* This paper is a heavily revised version of a chapter in my PhD thesis. The revision is based on comments from my dissertation committee, Thore Johnsen, Finn Kydland, Morten Ravn, and Erling Steigum, to whom I am grateful. I thank Sven-Arne Persson and participants at FIBE 2006 for comments to the current version. Remaining errors are solely my responsibility.

** Norwegian University of Science and Technology (NTNU) and Norges Bank. Address: NTNU, Department of Economics, NO-7491 Trondheim, Norway. Email: egilm@svt.ntnu.no

1. Introduction

Investors commonly delegate the management of their resources to several professional managers. Examples include mutual fund investors that spread their holdings between different fund managers, venture capitalists that allocate capital to many entrepreneurs, and direct investments in firms governed by different CEOs. The allocation decision faced by an investor in such cases can be described as a portfolio choice where the return on available assets is, in part, determined by the actions of different agents. Indeed, one may argue that the *purpose* of delegated investment management is to realize a rate of return that managers, but not the investor, can possibly achieve. As a large literature on optimal contracts makes clear,¹ the actions of managers under such investment arrangements are influenced by the compensation schemes offered to them. In a delegated investment setting, a rational investor must thus choose the optimal managerial contracts, in addition to the optimal portfolios.

Despite the frequent occurrence of delegated investment arrangements, standard models of portfolio choice usually assume that investors manage their own wealth,² making them unsuitable for analyzing the type of investment decisions discussed here. Models of optimal contracts have been applied to portfolio choice (Sung, 1995; Dybvig *et al.*, 2001; Ou-Yang, 2003; Westerfield, 2005), but these applications analyze how the portfolio choice of *agents* (i.e., managers) responds to contractual incentives; they do not discuss how an investor's portfolio selection interacts with optimal contracts.³

¹ See, e.g., Holmstrom and Milgrom (1987), Schättler and Sung (1993, 1997), Sung (1995, 2005), Müller (1998), Ou-Yang (2003), and Westerfield (2005).

² See Campbell and Viceira (2002) for a synthesis of modern theories of portfolio choice.

³ The contract literature referred to here discusses models with one principal and one agent, which obviously precludes an analysis of how an investor optimally should allocate wealth among several managers.

This paper investigates the portfolio decisions of an investor that can allocate her wealth between two projects. The mean rate of return from each project is determined by the actions (effort) of the manager in charge of it. These actions may or may not be observed by the investor, and are affected by the contract between investor and the managers. The model draws on the analysis of dynamic principal–agent problems by Holmstrom and Milgrom (1987), Schättler and Sung (1993, 1997), Sung (1995), and Müller (1998). That is, we explore a continuous-time model where output follows a Brownian motion and both the principal (the investor) and the agents (the managers) have constant absolute risk aversion, defined over terminal wealth⁴. This is a natural point of departure as, unlike the static principal–agent problem, the continuous-time version admits relatively easily interpreted solutions.

Upon presenting the model in Section 2, we analyze the first-best case of observable and enforceable actions in Section 3. We show that, even in this case, the risk–return trade-off involved in the investor’s portfolio decision is quite different from that in a standard portfolio choice model. First, the investor’s risk aversion is effectively lower than her CARA-coefficient, because some of the terminal wealth risk is carried by the managers, according to the optimal contracts. Second, there are important feedback effects, working through optimal contracts and efforts, that make expected terminal wealth nonlinear in initial wealth allocation. Depending on the shape of managers’ cost functions, the investor may have incentives to choose a highly “nondiversified” portfolio, even if the projects are completely symmetric. A general insight is that, in a principal–agent setting, the rate of return on a given investment is *endogenous* to the level of

⁴ Williams (2004) and Sannikov (2004) study moral hazard models where the (single) agent can consume continuously. Williams’s model is very general, but in most cases the solution can only be characterized and not solved explicitly. Sannikov assumes a risk neutral principal, making that model less suitable for studying an investor’s portfolio choice.

investment in that project; standard portfolio choice models treat the rate of return (and risk) on available assets as exogenous.

In Section 4, we investigate the case where there are hidden actions in one of the projects.⁵ This introduces an asymmetry where, *a priori*, one perhaps would expect a tilt in the portfolio away from the hidden action project. We do indeed show that the expected return in this project is generally lower than in the first-best. However, the relevant risk in the project, as seen from the investor's point of view, may be *lower* with moral hazard. This is because the investor uses the output from the project where actions can be observed as a signal to be included in the contract with the other manager (Holmstrom, 1982). The optimal contract then implies that the manager with hidden actions must carry more of the idiosyncratic risk associated with his project. Consequently, the minimum-variance allocation of the investor's wealth is in most cases tilted towards the hidden action project. Even if expected returns generally would be lower in this project, there are therefore plausible cases where the investor would in fact invest *more* in it than if she had full information.

Section 5 concludes the paper, and proofs are in the appendix.

2. The principal-agent framework

⁵ Think of a fund that employs some managers "in-house", with the possibility of internal monitoring, and delegates some of its assets to external managers, where the possibilities of monitoring are limited. Another example can be taken from the literature on the "home-bias puzzle" (see Lewis, 1999, for a survey), where it is sometimes argued that investors hold the lion's share of their portfolios domestically because there is asymmetric information across countries. A possible interpretation of the setup in Section 4 is that of an investor who can make direct investments in imperfectly correlated domestic and foreign projects, and where it is easier to monitor the actions of a domestic manager.

We investigate the principal–agent relationship on the time interval $[0,1]$. At time 0, the principal (the investor) decides how to allocate initial wealth $W(0) = W_0$ between two projects, A and B . The investment decision is assumed to be irreversible; the allocation is fixed until time 1.⁶ The output from the projects is publicly observable and governed by the processes

$$dX_i(t) = f(u_i(t), X_{i0}, t)dt + \sigma_i X_{i0} dz_i(t), \quad i = A, B, \quad (1)$$

where X_{i0} is the amount invested project i , $X_{A0} + X_{B0} = W_0$. Furthermore, σ_i is a diffusion parameter and dz_i is a standard Wiener process that represents a project-specific shock. The instantaneous correlation coefficient ρ of these shocks is obtained from $dz_A dz_B = \rho dt$, $\rho \in [-1,1]$. Throughout, we will assume that the production function f is given by

$$f(u_i, X_{i0}, t) = u_i(t)X_{i0}, \quad i = A, B. \quad (2)$$

The variable u_i (later referred to as “effort”) is controlled by the manager in charge of project i (manager i), and may or may not be observed by the investor. In any case, the expected rate of return on invested resources in project i is controlled by manager i .

This setup implies that the investor accumulates wealth according to

$$dW(t) = [\omega(u_A(t) - u_B(t)) + u_B(t)]W_0 dt + W_0 \mathbf{w}' \Sigma d\mathbf{B}(t), \quad (3)$$

where

$$\Sigma = \begin{bmatrix} \sigma_A & 0 \\ \sigma_B \rho & \sigma_B \sqrt{1 - \rho^2} \end{bmatrix},$$

ω is the fraction of initial wealth invested in project A , $\mathbf{w}' \equiv [\omega \quad 1 - \omega]$, $d\mathbf{B} \equiv [dz_A \quad dh]'$,

and dh is a standard Wiener process with the property that it is independent of dz_A .

⁶ This assumption is imposed to obtain tractability, as allowing for continuous reallocation would introduce time-dependent drifts in the processes for X_A and X_B . Schättler and Sung (1997) show that introducing time-dependent drifts of the Brownian motions would destroy the result that sharing rules are linear in output, and therefore the tractability of the model.

At time 0, the investor and the managers individually agree on sharing rules specifying payment from the investor to the managers at time 1. The sharing rules specify salaries S_A and S_B for manager A and B , respectively, and these are stochastic via dependence on the outcome of the stochastic process for W . The managers' control variables, $u_A \geq 0$ and $u_B \geq 0$, can be revised continuously during the time interval $[0,1]$ and may depend on the history of W in $[0,t]$, but not on the future $(t,1]$.

The managers incur costs of effort, assumed, for simplicity, to be quadratic in the level of effort

$$c(u_i(t), X_{i0}, t) = \frac{1}{2} u_i^2(t) X_{i0}^\theta, \quad i = A, B,$$

where $\theta \geq 0$. The assumptions of linear output and convex costs in effort are sufficient for well-defined solutions to both the first- and the second-best problems discussed below; see Theorem 4.2 in Schättler and Sung (1993). Observe that cost functions are symmetric: for a given allocation of resources, the cost of effort is equal for the two managers; for given effort, the cost of managing wealth is equal.

Finally, both the investor and the managers have exponential utility functions.

The investor's constant coefficient of absolute risk aversion is R while the two managers are equally risk averse with a CARA-coefficient r .

3. Full information

We shall first characterize the optimal sharing rules, effort levels, and resource allocation in the first-best setting; that is, when both managers' controls are observable and can be enforced at no cost. This analysis demonstrates that the trade-offs involved in the diversification decision are quite different from a standard portfolio choice model, even

when the investor can observe the actions of the managers. In addition, the analysis provides a benchmark for the case of asymmetric information considered below.

3.1. General solution

Ignoring time discounting, the investor's first-best problem at time 0 is

$$\max_{\{u_A, u_B\}, S_A, S_B, \omega} E\left[-\exp\{-R(W(1) - S_A - S_B)\}\right], \quad (4)$$

subject to (4) and to the managers' participation constraints:

$$E\left\{-\exp\left[-r\left(S_i - \int_0^1 c(u_i, X_{i0}, t)dt\right)\right]\right\} \geq -\exp\{-rU_0\}, \quad i = A, B, \quad (5)$$

where U_0 is the managers' certainty equivalent at time 0, assumed to be identical for the two managers. The solution to this problem is summarized in the first result.

Proposition 1: *Under full information, the effort levels are constant, unique, and determined by the equality between the marginal productivity and the marginal cost of effort:*

$$u_i = X_{i0}^{1-\theta}, \quad i = A, B. \quad (6)$$

Moreover, the salaries of the two managers are linear in combined output:

$$S_i = K_i + \frac{R}{r + 2R} W(1), \quad (7)$$

where

$$K_i \equiv \frac{1}{r + 2R} \left[\ln\left(\frac{\lambda r}{R}\right) - R W_0 + \frac{1}{2} r u_i^2 X_{i0}^\theta + \frac{1}{2} R (u_i^2 X_{i0}^\theta - u_j^2 X_{j0}^\theta) \right], \quad i = A, B; i \neq j,$$

is a constant. Finally, the investor allocates a fraction

$$\omega = \left[\frac{u_A - u_B}{W_0(\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB})R} \right] \frac{r + 2R}{r} + \left[\frac{\frac{1}{2}\theta(u_B^2 X_{B0}^{\theta-1} - u_A^2 X_{A0}^{\theta-1})}{W_0(\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB})R} \right] \frac{r + 2R}{r} + \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} \quad (8)$$

of initial wealth to project A.

Proof: See the appendix.

The effort levels u_A and u_B are constant because ω is constant and W_0 is given. We note from (6) that the first-best effort levels generally depend on the investor's allocation decision. Managers' effort is increasing (decreasing) with resources under management if θ is smaller (greater) than one. This is because the managers' marginal productivity of effort increases faster (more slowly) in allocated resources than do marginal costs when $\theta < (>) 1$, giving incentives to increase (decrease) effort. Note also that, given the optimal effort levels, the production functions in (2) are increasing in X_0 if $\theta < 2$, and have increasing returns in X_0 if $\theta < 1$.

The optimal sharing rules given in (7) are similar to the corresponding rule in the one-agent model of Müller (1998). One difference is that the coefficient before $W(1)$ gives more weight to the principal's risk aversion, as she now shares the final output with two agents. Moreover, the first-best sharing rules imply full risk sharing between the two managers. They receive a fixed share of total output, independent of the relative output from the project of which they are in charge. The constant amounts paid to the two managers differ to the extent that their effort costs differ. The agent with the highest effort cost receives the biggest constant amount.

The demand function (8) describes the investor's optimal diversification with full information. The first term on the right-hand side represents demand arising from potentially higher return on one of the projects. Relative to a standard portfolio selection problem (see Ch. 2 in Campbell and Viceira, 2002, for a recent exposition), this demand is adjusted by a factor $(r + 2R)/r$: the inverse of the share of final wealth retained by the investor according to the optimal contracts. Any given expected return difference has a

bigger impact on the allocation decision, the higher the share of final wealth retained by the investor. The second term is new, compared to standard portfolio choice problems. It reflects possible differences in managers' marginal wealth management costs. If, for a given allocation, manager A 's marginal cost is higher than B 's, it will contribute to lowering the allocation to project A . This is due to the properties of the constant amounts specified in the optimal contracts, as discussed above. The last term in (8) is the portfolio share that minimizes the variance of time 1 wealth. If marginal expected returns and marginal costs are equal across projects, the minimum-variance portfolio is optimal.

3.2. A mean-variance interpretation

Proposition 1 does not give a closed-form solution to the first-best problem; this is possible for a few values of θ only (to be discussed below). The proposition does however characterize the trade-offs involved in the investor's allocation decision.

To see how, we first recall that the investor derives utility from net final wealth.

We can use Proposition 1 to write net final wealth as

$$W(1) - S_A - S_B = \frac{r}{r + 2R} \left[W(1) - \frac{1}{2} (u_A^2 X_{A0}^\theta + u_B^2 X_{B0}^\theta) \right]. \quad (9)$$

Next, observe that constant first-best effort levels imply that the wealth process (3) follows an arithmetic Brownian motion. It follows that the gross return on wealth over the time interval $[0,1]$ is normal with mean $1 + \omega u_A + (1-\omega)u_B$ and standard deviation $\mathbf{w}'\Sigma$. With normally distributed returns, maximizing (4) is equivalent to maximizing the mean-variance utility function

$$E[W(1) - S_A - S_B] - \frac{1}{2} R \text{var}[W(1) - S_A - S_B].$$

Taking the expectation in (9) and using (6), we have

$$E[W(1) - S_A - S_B] = \frac{r}{r + 2R} \left[W_0 + \frac{1}{2} (\omega W_0)^{2-\theta} + \frac{1}{2} ((1-\omega)W_0)^{2-\theta} \right], \quad (10)$$

while the variance is

$$\text{var}[W(1) - S_A - S_B] = \left(\frac{r}{r+2R}\right)^2 W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w}. \quad (11)$$

We can thus write the investor's mean-variance utility as

$$\left[W_0 + \frac{1}{2}(\omega W_0)^{2-\theta} + \frac{1}{2}((1-\omega)W_0)^{2-\theta}\right] - \frac{rR}{r+2R} W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w}. \quad (12)$$

This is an indirect utility function, characterizing the investor's risk-return trade-off after incorporating the first-best effort levels and contracts.

The mean-variance function in (12) highlights two differences from a standard portfolio choice model. First, the effective risk aversion is $rR/(r+2R) < R$. Some of the final wealth risk is carried by the managers, implying less aversion to a given wealth variance of the investor. The investor's risk tolerance, or risk-bearing capacity, is $(r+2R)/rR$, compared to $1/R$ in the standard model. The difference in risk tolerance between the two models is thus $2/r$; the less risk averse the managers, the higher is the investor's tolerance of final wealth variance compared to the standard model.

Second, expected net final wealth is generally a nonlinear function of initial wealth allocation (it is linear in the standard model). This nonlinearity occurs because of feedback effects from wealth allocation to expected net final wealth; allocation of wealth determines effort and salaries, which in turn determines expected net final wealth.

Let the function $F(\omega)$ denote expected net final wealth, as given in the square brackets of (12), defined over the interval $\omega \in [0,1]$.⁷ The properties of $F(\omega)$ depend on the size of θ , as does accordingly the investor's portfolio choice.

Suppose first that $\theta \in (0,1)$. We can then establish that $F(\omega)$ is strictly convex in $\omega \in [0,1]$ and that it is symmetric around $\omega = \frac{1}{2}$. This implies that expected net final

⁷ From (6), short positions in any of the projects would violate the assumptions of nonnegative effort levels.

wealth is maximized by investing all initial wealth in *either* A or B . Underlying this property is that, by (6), effort levels are increasing in X_0 and, by (2), there are increasing returns in the projects. The consequence for portfolio choice depends on the relative size of σ_A and σ_B . If $\sigma_A < \sigma_B$, the minimum-variance position requires more than 50% of initial wealth in project A . In addition, however, there is now a net expected portfolio return to be gained by investing even more in A . The investor is generally willing to trade some final wealth variance against this increase in expected returns and the optimal portfolio is thus in the interval between the minimum-variance portfolio and $\omega = 1$. Analogously, $\sigma_A > \sigma_B$ implies that the optimal portfolio share ω is in the interval between 0 and the minimum-variance position. In summary, when $\theta \in [0,1)$ the endogenous responses of effort levels imply that the optimal portfolio is tilted beyond the composition that minimizes final wealth variance.⁸

The next case is $\theta = 1$. Equation (6) now implies that effort levels are independent of wealth allocation, implying constant expected returns in the two projects. This corresponds to a standard portfolio choice model. Moreover, as $u_A = u_B = 1$, expected returns are equal in A and B , implying that the minimum-variance portfolio is optimal.

Now turn to the case where $\theta \in (1,2)$. We can show that $F(\omega)$ is symmetric around $\omega = \frac{1}{2}$ in this case as well, but it is now strictly concave over $\omega \in [0,1]$. Expected net final wealth is therefore maximized by investing equal amounts in A and B . This is because effort levels are decreasing in X_0 , giving decreasing returns of the production functions in (2). Hence, any reallocation towards, for example, A from $\omega = \frac{1}{2}$ gives a

⁸ If $\sigma_A = \sigma_B$ and $\theta \in [0,1)$, the model gives ambiguous predictions as to whether the portfolio would be tilted towards A or B .

smaller gain in terms of increased expected output from A than the loss of expected output from B . The optimal portfolio is between $\omega = \frac{1}{2}$ and the minimum-variance portfolio, with the tilt being towards A (B) if $\sigma_A < (>) \sigma_B$. The minimum-variance portfolio is the optimal portfolio if $\sigma_A = \sigma_B$.

With $\theta = 2$, $F(\omega)$ is a constant. Expected net final wealth is thus independent of the allocation decision, and the optimal portfolio is the minimum-variance portfolio.

Finally, consider the case where $\theta > 2$. Equations (2) and (6) now imply that output from both projects falls with invested resources, making this case perhaps less interesting. Anyway, we can show that expected net final wealth is again maximized by investing all resources in one project only. The optimal portfolio is accordingly located in the same intervals as in the case with $\theta \in [0,1)$.

3.3. Closed form solutions in three special cases

If managers' costs are either independent of wealth ($\theta = 0$), linear in wealth ($\theta = 1$), or quadratic in wealth ($\theta = 2$), we can solve the investor's portfolio in closed form. The latter two cases were discussed above; they both imply that the minimum-variance portfolio is optimal.

With $\theta = 0$, we can use (8) to show that the optimal share of initial wealth invested in project A is

$$\omega = \frac{rR(\sigma_B^2 - \sigma_{AB}) - (r + 2R)}{rR(\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}) - 2(r + 2R)}. \quad (13)$$

When $\sigma_A = \sigma_B$, the optimal and minimum-variance portfolios coincide at $\omega = \frac{1}{2}$. If $\sigma_A \neq \sigma_B$, the optimal portfolio is tilted beyond what is implied by the minimum-variance

portfolio. The intuition was explained in the last subsection: as $\theta < 1$, the investor can increase expected returns by investing in one project only.

4. Diversification with asymmetric information

We now turn to the case where the investor cannot observe the actions of manager B .

The investor faces asymmetric information along two dimensions: she knows less about the actions of manager B than does the manager himself, and she knows more about the actions of manager A than of B 's. We are primarily interested in the consequences of these asymmetries for diversification and portfolio choice. In particular, does less information necessarily mean less investment?

4.1. The second-best problem

With moral hazard in project B , the investor faces an additional constraint:

$$u_B \in \arg \max_{u_B} E \left[-\exp \left\{ -r \left(S_B - \int_0^1 c(u_B, X_{B0}, t) dt \right) \right\} \right].$$

This is the familiar incentive compatibility constraint, which says that manager B chooses the u_B that is in his best interest. We follow Schättler and Sung (1993), and use the so-called first-order approach to solve the investor's problem. In this approach, the incentive compatibility constraint in the principal's problem is relaxed to the first-order necessary condition for optimality in the agent's problem.

As the investor has full information on the actions of manager A , and as the managers do not control variances of the Brownian processes, the output from project A serves as a signal to be included in the contract between the investor and manager B . In

particular, following Holmstrom and Milgrom (1987, pp. 324–25), we assume that the salary to manager B is conditioned on the aggregate performance index

$$\begin{aligned} Y(t) &= X_B(t) - \frac{\text{cov}(X_A(t), X_B(t))}{\text{var}(X_A(t))} X_A(t) \\ &= X_B(t) - \frac{\rho\sigma_B X_{B0}}{\sigma_A X_{A0}} X_A(t). \end{aligned}$$

Using (1), (2), and (3), we can show that this index evolves according to:

$$dY(t) = X_{B0} \left[u_B(t) - \frac{\rho\sigma_B}{\sigma_A} u_A(t) \right] dt + X_{B0} \mathbf{m}' \Sigma d\mathbf{B}(t), \quad (14)$$

where $\mathbf{m}' \equiv \left[-\frac{\rho\sigma_B}{\sigma_A} \quad 1 \right]$.

Now, the problem of manager B is

$$\max_{u_B} E \left[-\exp \left\{ -r \left(S_B - \int_0^1 c(u_B, X_{B0}, t) dt \right) \right\} \right],$$

subject to (14), taking as given the actions of manager A and the allocation decision of the investor. Applying the representation given in Schättler and Sung (1993), we can show that the solution to this problem gives the optimal sharing rule:

$$S_B = U_0 + \frac{1}{2} \int_0^1 u_B^2 X_{B0}^0 dt + \int_0^1 u_B X_{B0}^0 \mathbf{m}' \Sigma d\mathbf{B}(t) + \frac{1}{2} r \int_0^1 (u_B X_{B0}^0)^2 \mathbf{m}' \Sigma \Sigma' \mathbf{m} dt. \quad (15)$$

The first two terms on the right-hand side of (15) provide manager B with his certainty equivalent plus compensation for the cost he actually incurs. The next term is the compensation error, arising because the investor's compensation is based on realized outcome rather than the manager's actual effort. Finally, to compensate manager B for the risk he carries, a risk-premium is paid, given by the last term in (15).

The investor's relaxed problem can then be written as

$$\max_{u_A, u_B, S_A, \omega} E \left[-\exp \left\{ -R(W(1) - S_A - S_B) \right\} \right],$$

subject to (3), (15), and manager A 's participation constraint. The solution is summarized in the second proposition.

Proposition 2: *Suppose that manager B's effort cannot be observed. Manager A's optimal effort level is still given by condition (6). Manager B's optimal effort is constant and fulfills:*

$$u_B = X_{B0}^{1-\theta} g(X_{B0}), \quad (16)$$

where

$$g(X_{B0}) \equiv \frac{1 + \frac{rR}{r+R} X_{B0}^\theta \sigma_B^2 (1-\rho^2)}{1 + \frac{r(r+2R)}{r+R} X_{B0}^\theta \sigma_B^2 (1-\rho^2)} \leq 1. \quad (17)$$

The optimal salary contract with manager B is linear in output from each of the projects:

$$S_B = \kappa_B + g(X_{B0}) \left(X_B(1) - \frac{\rho \sigma_B X_{B0}}{\sigma_A X_{A0}} X_A(1) \right), \quad (18)$$

where

$$\kappa_B \equiv U_0 + \frac{1}{2} u_B^2 X_{B0}^\theta + \frac{1}{2} \sigma_B^2 (1-\rho^2) [g(X_{B0})]^2 X_{B0}^2 - g(X_{B0}) E[Y(1)]$$

is a constant. The salary of manager A is given by

$$S_A = \kappa_A + \frac{R}{r+R} \left[\left(1 + g(X_{B0}) \frac{\rho \sigma_B X_{B0}}{\sigma_A X_{A0}} \right) X_A(1) + (1 - g(X_{B0})) X_B(1) \right], \quad (19)$$

with

$$\kappa_A \equiv \frac{1}{r+R} \left[\ln \left(\frac{\lambda_A r}{R} \right) - R W_0 + \frac{1}{2} r u_A^2 X_{A0}^\theta - R \kappa_B \right]$$

being another constant. Finally, the investor allocates a fraction

$$\omega = \left(\frac{u_A - u_B - \frac{1}{2} \theta (u_A^2 X_{A0}^{\theta-1} - u_B^2 X_{B0}^{\theta-1})}{W_0 [\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}] R} \right) \frac{r+R}{r} + \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} - \left(\frac{\sigma_B^2 (1-\rho^2) u_B X_{B0}^\theta (1 + \theta - \frac{r+2R}{R} \theta u_B X_{B0}^{\theta-1})}{W_0 [\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}]} \right), \quad (20)$$

of initial wealth to project A.

Proof: See the appendix.

As expected, manager B 's effort is (weakly) lower than A 's for given investments in the respective projects; see equations (16) and (6). The variable $g(X_{B0})$, defined in (17), can be interpreted as a wedge between the first- and second-best efforts, conditional on the amount invested in B . We can show that $g'(X_{B0}) \leq 0$, where the equality holds if $\theta = 0$ or $\rho = -1$. Except in these two special cases, manager B 's effort level will be lower *relative* to the first-best the more the investor allocates to project B . Whether more investment yields higher effort in *absolute terms* depends on the sign of

$$\frac{du_B}{dX_{B0}} = (1 - \theta)X_{B0}^{-\theta}g + X_{B0}^{1-\theta}g'(X_{B0}). \quad (21)$$

As long as $g'(X_{B0}) < 0$, we now need $\theta \ll 1$ for manager B 's effort to respond positively on investment. Compared to the first-best, there is thus smaller range of θ 's for which effort increases with investment.

The sharing rule (18) resembles the relative evaluation scheme discussed by Holmstrom and Milgrom (1987, p. 324), but here the shares of output accruing to the investor and the managers are endogenous to the allocation of initial wealth. Note that (18) can be written as

$$S_B = \tilde{\kappa}_B + g(X_{B0})(Y(1) - E[Y(1)]),$$

where $\tilde{\kappa}_B = \kappa_B + gE[Y(1)]$. Manager B is given a constant amount plus a share g of the *surprise* in the aggregate performance index Y . The third term in the definition of κ_B is the risk-premium that is optimally paid to manager B . Note that it depends on the idiosyncratic risk $\sigma_B^2(1 - \rho^2)$ associated with project B . Finally, (18) and (19) show that risk sharing between the two managers is generally imperfect when there is moral hazard.

Turning to portfolio choice, we see that the upper line in (20) bears close resemblance to equation (8). The term in the lower line, however, is new and conceptually different. It occurs because initial wealth allocation affects the risk-premium

and the expected value of the performance index (the last two terms in the definition of κ_B). In turn, this feeds back to the expected value and variability of investors' net final wealth, and therefore their portfolio decisions. To give a more thorough interpretation of portfolio choice, it is again helpful to recast the portfolio decision within a mean-variance setting.

4.2. The risk-return trade-off with moral hazard

We can use (18), (19), and the performance index Y to express the investor's net final wealth as (we suppress the argument in $g(\cdot)$ to simplify notation):

$$\begin{aligned} & W(1) - S_A - S_B \\ &= \frac{r}{r+R} \left[W(1) + g(E[Y(1)] - Y(1)) - \frac{1}{2} (u_A^2 X_{A0}^\theta + u_B^2 X_{B0}^\theta) - \frac{1}{2} r g^2 X_{B0}^2 \sigma_B^2 (1 - \rho^2) \right] + tip. \end{aligned} \quad (22)$$

Now, as both managers' efforts are constant also in the second-best, the processes (3) and (14) both follow arithmetic Brownian motions. It follows from equation (22) that net final wealth is normally distributed over the time interval $[0,1]$.

Maximizing (4) is thus again equivalent to maximize

$$E[W(1) - S_A - S_B] - \frac{1}{2} R \text{var}[W(1) - S_A - S_B].$$

Taking expectations in (22) and substituting from (6) and (17), we find:

$$\begin{aligned} & E[W(1) - S_A - S_B] \\ &= \frac{r}{r+R} \left[W_0 + \frac{1}{2} (\omega W_0)^{2-\theta} + g(1 - \frac{1}{2} g) ((1 - \omega) W_0)^{2-\theta} - \frac{1}{2} r \sigma_B^2 (1 - \rho^2) g^2 ((1 - \omega) W_0)^2 \right] \\ &= \frac{r}{r+R} \left[F(\omega) - (\frac{1}{2} - g(1 - \frac{1}{2} g)) ((1 - \omega) W_0)^{2-\theta} - \frac{1}{2} r \sigma_B^2 (1 - \rho^2) g^2 ((1 - \omega) W_0)^2 \right], \end{aligned} \quad (23)$$

where the function $F(\omega)$ was defined and analyzed in Section 3. Comparing (10) and (23), we see that the link between the initial allocation and expected net final wealth is quite different in the two cases. The last two terms in the lower line of (23) are new relative to the first-best. The first of these shows the loss in expected output net of management

costs in project B , compared to the first-best. We can show that this loss is strictly increasing over $\omega \in [0,1]$, as long as $\theta \leq 2$. When $\theta > 2$, the loss is not necessarily monotonic over $\omega \in [0,1]$, but still reaches its minimum of zero at $\omega = 1$. The final term in (23) is the risk-premium paid to manager B . It enters with a minus in front, so it contributes to lowering the investor's net expected wealth. The risk-premium also has a minimum of zero at $\omega = 1$, and can be shown to be strictly decreasing over $\omega \in [0,1]$ when $\theta \leq 1$.

These observations allow us to deduce the following on wealth allocation and expected net final wealth in the second-best. Starting with the cases where $\theta \notin [1,2]$, we recall from Section 3 that $F(\omega)$ then has maxima at 0 and 1. However, as the last two terms in the lower line of (23) both have a unique maximum at $\omega = 1$, expected net final wealth is greatest when investing in project A only. The cases with $\theta = \{1,2\}$ are analogues; $F(\omega)$ is independent of wealth allocation, and hence expected net final wealth with moral hazard is maximized by setting $\omega = 1$. The last possibility is $\theta \in (1,2)$, where we recall that $F(\omega)$ has its maximum at $\omega = \frac{1}{2}$. The last term in the upper line of (23) gives an incentive to invest more in project A . However, as the risk-premium to manager B is possibly non-monotonic in ω for these values of θ , the overall link between initial wealth allocation and expected net final wealth is theoretically ambiguous.⁹

Turning to the risk of net final wealth, we observe that (22) implies

$$\text{var}[W(1) - S_A - S_B] = \left(\frac{r}{r+R}\right)^2 \text{var}[W(1) - gY(1)].$$

⁹ Numerical calculations by the author indicate that, for plausible parameters, expected net final wealth is maximized for $\omega > \frac{1}{2}$ also when $\theta \in (1,2)$.

The optimal contract with manager B leaves the investor exposed to the risk of *gross* final wealth less a (endogenous) share g of the aggregate performance index Y . This risk exposition is in turn shared with manager A (the investor carries a share $r/(r+R)$); hence the above expression. Using equation (14), we have

$$\text{var}[W(1) - S_A - S_B] = \left(\frac{r}{r+R}\right)^2 \left[W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w} - g(2-g)((1-\omega)W_0)^2 \sigma_B^2 (1-\rho^2) \right].$$

An implication of this equation is that, unlike with full information, the portfolios that minimize *gross* and *net* final wealth variance, respectively, are not equal. Interestingly, we can show that the net final wealth minimizing share invested in project A is smaller than $(\sigma_B^2 - \sigma_{AB})/(\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB})$ unless $\theta \gg 2$. The most plausible cases of the model thus imply that the relevant minimum-variance position is tilted towards project B when there is moral hazard.

The implications for portfolio choice of the above discussion are straightforward. Relative to the first-best, expected returns can generally be increased by tilting the portfolio towards project A . On the other hand, the minimum-variance position would in most cases be tilted towards project B . The combined effect on portfolio choice is therefore generally ambiguous. Interestingly, we cannot rule out that the investor should optimally *increase allocation to the project with moral hazard* compared with the case of full information.

We close this section by illustrating the latter point using an example with $\theta = 0$, the only case where the portfolio rule can be expressed in closed form. For simplicity, we set $\sigma_A = \sigma_B = \sigma$ (implying an optimal first-best rule $\omega = 1/2$; see the discussion in Section 3.3), and assume that $\rho = 0$. When $\theta = 0$, the wedge g is a constant given by

$$g = \frac{r + R + rR\sigma^2}{r + R + r(r + R)R\sigma^2}.$$

Substituting into (20), we can show that the optimal share invested in A is

$$\omega = \frac{rR\sigma^2 (r^2\sigma^2 - r - R + rR\sigma^2) - (r + R)^2}{2rR\sigma^2 (r^2\sigma^2 - \frac{3}{2}r - R + \frac{3}{2}rR\sigma^2) - 2(r + R)^2 - r^3\sigma^2}.$$

It is straightforward to show that this share can be less than half for plausible parameter values.

5. Conclusions

We have studied the resource allocation of investors that delegate the management of their wealth to two different managers. Managers' effort levels determine the expected return from the projects that they govern, but these effort levels are affected by the contracts offered by investors. We show that even when managers' actions are observable and enforceable, the investors' diversification decision involves trade-offs other than a standard portfolio problem. This is because, in a principal-agent setting, expected returns are endogenous to the allocation of initial wealth. Depending on the shape of managers' cost functions, the expected final wealth net of managerial compensation can exhibit both increasing and decreasing returns to invested wealth. In the former case, optimal portfolio holdings may be highly undiversified despite the symmetry between projects in the underlying model.

If managers' actions cannot be observed, additional mechanisms come into play. We explore the case with moral hazard in one of the projects. Relative to the first-best, expected returns net of managerial costs can generally be increased by tilting the portfolio away from the moral hazard project. However, the optimal contract with the manager of this project transfers more of the idiosyncratic risk associated with the

project to the manager. As a consequence, the minimum-variance position of the investors' net final wealth would in most cases be tilted *towards* the moral hazard project. There are plausible cases of the model where the allocation to the project with hidden actions is higher than in the first-best.

Appendix

A.1. Proof of Proposition 1

We start by deriving the optimal sharing rules in terms of the optimal controls u_A and u_B , following Müller (1998). Net compensation to manager i is $y_i = S_i - \frac{1}{2} X_{i0}^\theta \int_0^1 u_i^2(t) dt$. Then, integrating (3) and inserting the result in (4) imply that the investor's problem can be expressed as

$$\max_{\{u_A, u_B\}, y_A, y_B, \omega} E \left[-\exp \left\{ -R(W_0 + W_0 \mathbf{w}' \Sigma (\mathbf{B}_1 - \mathbf{B}_0)) - y_A - y_B \right. \right. \\ \left. \left. + \int_0^1 \left[u_A(t) \omega W_0 - \frac{1}{2} u_A^2(t) (\omega W_0)^\theta + u_B(t) (1 - \omega) W_0 - \frac{1}{2} u_B^2(t) ((1 - \omega) W_0)^\theta \right] dt \right\} \right],$$

subject to (5). Pointwise maximization with respect to y_A and y_B gives the first-order conditions

$$y_i = \frac{1}{r + R} \ln \left(\frac{\lambda_i r}{R} \right) + \frac{R}{r + R} \left[W(1) - W_0 - \frac{1}{2} \int_0^1 \left(u_A^2(\omega W_0)^\theta + u_B^2((1 - \omega) W_0)^\theta \right) dt \right] - \frac{R}{r + R} y_j,$$

where $i, j = A, B$, $i \neq j$, and λ_i is the Lagrange-multiplier associated with (5). Solving these two equations for y_A and y_B and using the two participation constraints to demonstrate that $\lambda_A = \lambda_B$, we find

$$y_A = y_B = \frac{1}{r + 2R} \ln \left(\frac{\lambda_A r}{R} \right) \\ + \frac{R}{r + 2R} \left(W(1) - W_0 - \frac{1}{2} \int_0^1 \left[u_A^2(\omega W_0)^\theta + u_A^2((1 - \omega) W_0)^\theta \right] dt \right). \quad (\text{A.1})$$

The optimal sharing rules are $S_i = y + \int_0^1 c(u_i(t), X_{i0}) dt$, $i = A, B$, where y is given in

(A.1).

Next, we substitute the optimal salary functions into (4). The investor's problem can then be simplified to

$$\max_{\{u_A, u_B\}, \omega} E \left[-\exp \left\{ -a \left(W(1) - b - \frac{1}{2} \int_0^1 [u_A^2(\omega W_0)^\theta + u_B^2((1-\omega)W_0)^\theta] dt \right) \right\} \right]$$

subject to (3), where $a \equiv rR/(r+2R)$ and $b \equiv (r/2R)W_0 - (r/2)\ln(\lambda_A r/R)$ are constants.

Let $V(t, W(t))$ be the investor's value function, giving the optimal remaining utility at time t . By Lemma A1 in Sung (1995), the value function solving the above problem satisfies the following dynamic programming equation:

$$\begin{aligned} 0 \equiv & \frac{\partial V}{\partial t} + \max_{\{u_A, u_B\}, \omega} \left\{ \frac{\partial V}{\partial W} [u_A \omega W_0 + u_B (1-\omega)W_0] \right. \\ & \left. + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w} + \frac{1}{2} a V(t, W) [u_A^2(\omega W_0)^\theta + u_B^2((1-\omega)W_0)^\theta] \right\}, \end{aligned} \quad (\text{A.2})$$

with the terminal condition being $V(1, W(1)) = -\exp[-a(W(1)-b)]$. From (A.2), the first-order conditions with respect to $u_A(t)$, $u_B(t)$, and ω are

$$u_i h(X_{i0}) = \frac{-\partial V / \partial W}{a V(t, W(t))} X_{i0}, \quad i = A, B \quad (\text{A.3})$$

$$\begin{aligned} \omega W_0 = & -\frac{\partial V / \partial W}{\partial^2 V / \partial W^2} \left[\frac{u_A - u_B}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} \right] - \frac{a V(t, W)}{\partial^2 V / \partial W^2} \left[\frac{\frac{1}{2} \theta u_A^2(\omega W_0)^{\theta-1} - \frac{1}{2} \theta u_B^2((1-\omega)W_0)^{\theta-1}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} \right] \\ & + \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} W_0. \end{aligned} \quad (\text{A.4})$$

Finally, we conjecture that the value function has the form

$$\begin{aligned} V(t, W(t)) = & -\exp \left\{ -a \left[W(t) - b + (1-t)(f(u_A, \omega W_0) + f(u_B, (1-\omega)W_0)) \right. \right. \\ & \left. \left. - \frac{1}{2} a W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w} - c(u_A, \omega W_0) - c(u_B, (1-\omega)W_0) \right] \right\}. \end{aligned} \quad (\text{A.5})$$

Using (A.5) in (A.3) and (A.4), we obtain equations (6)–(8). Substituting (A.5) into (A.2) confirms that (A.5) solves the investor's dynamic problem.

A.2. Proof of Proposition 2

We can proceed as under full information to find the optimal sharing rule between the investor and manager A in terms of the optimal control u_A . The first-order condition with respect to S_A for the investor's relaxed problem is thus

$$S_A = \frac{1}{r+R} \ln\left(\frac{\lambda_A r}{R}\right) + \frac{R}{r+R} (W(1) - W_0 - S_B) + \frac{1}{2} \frac{r}{r+R} X_{A0}^\theta \int_0^1 u_A^2 dt. \quad (\text{A.6})$$

Given the optimal sharing rules in (15) and (A.6), the investor's (stochastic) net terminal wealth can be expressed as

$$W(1) - S_A - S_B = \frac{r}{r+R} \left[W(1) - \frac{1}{r} \ln\left(\frac{\lambda_A r}{R}\right) + \frac{R}{r} W_0 - U_0 - \frac{1}{2} \int_0^1 [u_A^2 X_{A0}^\theta + u_B^2 X_{B0}^\theta] dt \right. \\ \left. - \int_0^1 u_B X_{B0}^\theta \mathbf{m}' \Sigma d\mathbf{B}(t) - \frac{1}{2} r \int_0^1 (u_B X_{B0}^\theta)^2 \mathbf{m}' \Sigma \Sigma' \mathbf{m} dt \right].$$

It follows that the investor's problem can be reduced to

$$\max_{\{u_A, u_B\}, \omega} E \left[-\exp \left\{ -\alpha \left(W(1) - \beta - \int_0^1 u_B X_{B0}^\theta \mathbf{m}' \Sigma d\mathbf{B}(t) \right. \right. \right. \\ \left. \left. \left. - \frac{1}{2} \int_0^1 [u_A^2 X_{A0}^\theta + u_B^2 X_{B0}^\theta + r (u_B X_{B0}^\theta)^2 \mathbf{m}' \Sigma \Sigma' \mathbf{m}] dt \right) \right\} \right]. \quad (\text{A.7})$$

subject to (3), where $\alpha \equiv rR/(r+R)$ and $\beta \equiv (1/r) \ln(\lambda_A r/R) - (R/r) + U_0$ are constants.

The dynamic programming equation becomes

$$0 \equiv \frac{\partial V}{\partial t} + \max_{\{u_A, u_B\}, \omega} \left\{ \frac{\partial V}{\partial W} [u_A X_{A0} + u_B X_{B0} + \alpha (u_B X_{B0}^\theta) W_0 \mathbf{w}' \Sigma \Sigma' \mathbf{m}] \right. \\ \left. + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w} \right. \\ \left. + \frac{1}{2} \alpha V(t, W(t)) [u_A^2 X_{A0}^\theta + u_B^2 X_{B0}^\theta + (r + \alpha) (u_B X_{B0}^\theta)^2 \mathbf{m}' \Sigma \Sigma' \mathbf{m}] \right\}, \quad (\text{A.8})$$

with the terminal condition being $V(1, W(1)) = -\exp[-\alpha(W(1) - \beta)]$. The first-order conditions with respect to u_A , u_B , and ω , respectively, read:

$$u_A X_{A0}^\theta = \frac{-\partial V / \partial W}{\alpha V(\bullet)} X_{A0}, \quad (\text{A.9})$$

$$u_B X_{B0}^\theta = \frac{-\partial V / \partial W}{\alpha V(\bullet)} \left[\frac{X_{B0} + \alpha X_{B0}^\theta W_0 \mathbf{w}' \Sigma \Sigma' \mathbf{m}}{1 + (r + \alpha) X_{B0}^\theta \mathbf{m}' \Sigma \Sigma' \mathbf{m}} \right], \quad (\text{A.10})$$

$$\omega W_0 = -\frac{\partial V / \partial W}{\partial^2 V / \partial W^2} \left[\frac{u_A - u_B - \alpha(1 + \theta) u_B \sigma_B^2 (1 - \rho^2) X_{B0}^\theta}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} \right] \\ - \frac{\alpha V(t, W)}{\partial^2 V / \partial W^2} \left[\frac{\frac{1}{2} \theta (u_A^2 X_{A0}^{\theta-1} - u_B^2 X_{B0}^{\theta-1}) - \theta(r + \alpha) \sigma_B^2 (1 - \rho^2) u_B^2 X_{B0}^{2\theta-1}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} \right] + \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}} W_0. \quad (\text{A.11})$$

We use

$$V(t, W_t) = -\exp \left\{ -\alpha \left[W(t) - \beta + (1-t) \left(u_A X_{A0} + u_B X_{B0} + \alpha u_B X_{B0}^0 W_0 \mathbf{w}' \Sigma \Sigma' \mathbf{m} \right) \right. \right. \\ \left. \left. - \frac{1}{2} \alpha W_0^2 \mathbf{w}' \Sigma \Sigma' \mathbf{w} - \frac{1}{2} \left[u_A^2 X_{A0}^0 + u_A^2 X_{B0}^0 + (r + \alpha) (u_B X_{B0}^0)^2 \mathbf{m}' \Sigma \Sigma' \mathbf{m} \right] \right] \right\},$$

as a trial solution for the value function. Taking the appropriate derivatives and substituting into (A.9)–(A.11) gives (16) and (20). Equations (18) and (19) are obtained by combining (16) and (6) with (15) and (A.6), respectively. Substituting the trial solution into (A.8) confirms that it solves the dynamic programming equation.

References

- Campbell, John Y., and Luis M. Viceira (2002): *Strategic Asset Allocation – Portfolio Choice for Long-Term Investors*, Oxford University Press, Oxford.
- Dybvig, Philip H., Heber K. Farnsworth, and Jennifer N. Carpenter (2001): “Portfolio performance and agency”, working paper, Washington University, St. Louis.
- Holmstrom, Bengt (1982): “Moral hazard in teams”, *Bell Journal of Economics* 13, 324–340.
- Holmstrom, Bengt, and Paul Milgrom (1987): “Aggregation and linearity in the provision of intertemporal incentives”, *Econometrica* 55, 303–328.
- Lewis, Karen (1999): “Trying to explain home bias in equities and consumption”, *Journal of Economic Literature* 37, 571–608.
- Müller, Holger M. (1998): “The first-best sharing rule in the continuous-time principal-agent problem with exponential utility”, *Journal of Economic Theory* 79, 276–280.
- Ou-Yang, Hui (2003): “Optimal contracts in a continuous-time delegated portfolio management problem”, *Review of Financial Studies* 16, 173–208.
- Sannikov, Yuily (2004): “A continuous-time version of the principal-agent problem”, working paper, Stanford University.
- Schättler, Heinz, and Jaeyoung Sung (1993): “The first-order approach to the continuous-time principal-agent problem with exponential utility”, *Journal of Economic Theory* 61, 331–371.
- (1997): “On optimal sharing rules in discrete- and continuous-time principal-agent problems with exponential utility”, *Journal of Economic Dynamics & Control* 21, 551–574.
- Sung, Jaeyoung (1995): “Linearity with project selection and controllable diffusion rate in continuous-time principal-agent problems”, *RAND Journal of Economics* 26, 720–743.

————— (2005): “Optimal contracts under adverse selection and moral hazard: A continuous-time approach”, *Review of Financial Studies* 18, 1021–1073.

Westerfield, Mark M. (2005): “A solution to the principal-agent moral hazard problem in continuous time”, working paper, University of Southern California.

Williams, Noah (2004): “On dynamic principal-agent problems in continuous time”, working paper, Princeton University.