

OPTIMAL COMPENSATION FOR
FUND MANAGERS OF UNCERTAIN TYPE:
INFORMATIONAL ADVANTAGES OF BONUS SCHEMES¹

ALEXANDER STREMME

Stern School of Business, New York University

OCTOBER 1999

Address for Correspondence:

Alexander Stremme

Stern School of Business

New York University

44 West 4th St, K-MEC # 7-73

New York, NY 10012

phone: (212) 998-0432

fax: (212) 995-4221

e-mail: astremme@stern.nyu.edu

web: <http://www.stern.nyu.edu/~astremme/>

¹I would like to express my gratitude to David Backus, Jennifer Carpenter, Frank Diebold, Kose John, Anthony Lynch, Richard Payne, Matt Richardson, Raghu Sundaram, and David Webb for encouragement and many helpful suggestions, and seminar participants at New York University and the London School of Economics for stimulating discussions. All remaining errors are mine.

Abstract

Performance-sensitivity of compensation schemes for portfolio managers is well explained by classic principal-agent theory as a device to provide incentives for managers to exert effort or bear the cost of acquiring information. However, the majority of compensation packages observed in reality display in addition a fair amount of *convexity* in the form of performance-related bonus schemes. While convex contracts may be explained by principal-agent theory in some rather specific situations, they have been criticized, both by the financial press as well as the academic literature, on the grounds that they may lead to excessive risk-taking. In this paper, we show that convex compensation packages, though likely to be myopically not optimal, may serve as a device to extract information about the ex-ante uncertain type of portfolio managers. Optimal contracts are thus determined by the trade-off between maximizing short-run expected returns on one hand, and long-run informational benefits on the other. In a discrete-time model, combining dynamic principal-agent theory with the theory of learning by experimentation, we characterize optimal incentive schemes and optimal retention rules for fund managers, consistent with empirical observations.

JEL CLASSIFICATION: C61, D82, D83, G11

KEYWORDS: Fund Manager Compensation, Portfolio Choice,
Asymmetric Information, Learning by Experimentation

1 Introduction

Compensation contracts between investment funds or banks and their portfolio managers in recent years have become a focus of attention for both academic research as well as public interest. Principal-agent theory addresses the design of optimal compensation contracts in the context of the moral hazard or adverse selection problems that result from the non-verifiability of the actions taken by the agent or the unobservability of the agent's type. Compensation contracts serve as mechanisms to provide incentives for the agent to exert effort or to acquire costly information. Therefore, the compensation received by the agent must depend on some verifiable measure of performance which, in the case of portfolio management, is typically given by the return achieved on the managed portfolio.

However, the overwhelming majority of compensation packages observed in reality consist of a flat base salary plus an optional bonus which is paid only if the manager's performance exceeds some given benchmark. In other words, the compensation received by the manager is typically a (weakly) convex function of the performance measure, much like the pay-off of a call option, see for example Oyer (1998). In the existing principal-agent literature, there are essentially two situations in which optimal compensation contracts may turn out to be convex: First, if the agent's preferences exhibit a higher level of risk-aversion than the principal's, convexity may be required to align the agent's preferences with those of the principal. Second, if the agent incurs increasing marginal dis-utility of exerting effort, (or increasing marginal cost of acquiring information,) optimal contracts may be convex to provide appropriate incentives. For example, in a model in which the dis-utility of effort is given by a convex function, Stoughton (1993) shows that optimal contracts are in fact quadratic. However, in cases where the agent's actions affect both the mean *and* the variance

of outcomes, convex incentive schemes will most likely induce excessive risk-taking. For example, in a model in which the agent can choose among investment projects of different risk, Lambert (1986) shows that contracts designed to extract effort in fact lead to overinvestment in the risky project.

In this paper, we take a different approach to determining the structure of optimal compensation contracts for portfolio managers. Rather than viewing compensation schemes merely as devices to extract effort, we focus on their informational aspects. In doing so, we combine the principal-agent literature with the theory of learning by experimentation. More specifically, we show how convex compensation contracts may serve as devices to extract information about the agent's ex-ante uncertain type, *despite* the fact that such contracts may lead to excessive risk-taking and thus may not be optimal myopically.

Evidence for persistent differences in relative fund performance has been found, among others, by Goetzmann and Ibbotson (1994), Brown and Goetzmann (1995), and Gruber (1996). Moreover, Chevalier and Ellison (1998) find significant differences in performance even on the level of individual managers. In other words, certain portfolio managers seem to possess superior skill in selecting optimal portfolios, or have better access to private information. In our model, such differences in skill or information will be reflected in the manager's type. We postulate that the effect of type on performance will be stronger the higher the level of risk involved. The intuition is that private information obviously plays a stronger role when investing in high-risk stocks rather than, say, Treasury securities. One might argue that at the same time as higher risk increases the effect of the agent's type on performance, it also makes the information contained in observed performance more noisy. While this certainly creates an interesting trade-off for the principal, it will not be the focus of the analysis developed here. In this paper, we show how compensation schemes can be used

to make portfolio managers reveal their true type by inducing them to invest aggressively. Obviously, this informational benefit comes at a cost, since the portfolios chosen by managers under such incentive schemes may not be optimal in the short run. Thus, the long-run informational benefits of compensation schemes are traded off against the loss incurred by deviations from short-run optimality. This type of trade-off is the characterizing feature of the literature on “learning by experimentation”; see for example Aghion, Bolton, Harris, and Jullien (1991) or Keller and Rady (1999). It is worth noting that the results developed in this paper do not only apply to portfolio management. In any principal-agent situation in which agents’ types have an effect of performance that becomes more pronounced with the risk taken, the principal faces the same type of trade-off between the benefits and the cost of experimentation.

Casual observations made on the trading floors of large banks seem to suggest that these banks indeed offer extremely aggressive bonus schemes especially to more junior portfolio managers. While this may seem irrational at first, this paper offers the explanation that such bonus schemes may be used as screening devices to separate the “good” portfolio managers from the “bad” ones. Empirical evidence in support of this explanation is manifold: Both Benston (1985) and Khorana (1996) report a significantly negative relationship between performance and management turnover, which indicates that performance is indeed used as a screening device. Moreover, Khorana (1996) finds that return risk increases prior to the replacement of a manager. While most authors attribute this observation to “window dressing” or “herd behaviour” as studied by Lakonishok, Shleifer, Thaler, and Vishny (1991) and Scharfstein and Stein (1990), our paper reverses the direction of causality: It is not the impending threat of dismissal that causes managers to take on more risk, but instead the increase in risk makes dismissal more likely.

We develop a discrete-time, infinite-horizon model in which risk-neutral agents (portfolio managers) may be hired by a single, risk-averse principal to manage the principal's investment funds. Contracts between an agent and the principal last for one period and specify the compensation that the agent receives, contingent on the return achieved. Agents are of different types according to their ability to forecast returns, but types are not observable by the principal *ex ante*. Instead, the principal forms beliefs about the agent's type, conditional on observed returns. Based on these beliefs, the principal may decide to fire and replace the agent at the end of any period. In other words, management turnover is endogenously determined on the basis of performance. This adds to the existing literature on learning by experimentation, where the environment in which learning takes place is either fixed or changes exogenously. To keep this paper short, we restrict the analysis to the case in which the compensation scheme is fixed over time. The implications of allowing the principal to adjust the contract according to the evolution of beliefs is dealt with in a second, forthcoming paper. We will however mention the corresponding results whenever we believe them to be of interest. While hired by the principal, agents will choose a portfolio strategy such as to maximize expected future compensation, taking into account the possibility of being fired. The principal on the other hand chooses a compensation scheme and firing rule such as to maximize expected utility.

There have been numerous studies that focus on principal-agent interactions in dynamic environments, for example Radner (1985), Fudenberg, Holmstrom, and Milgrom (1986), or Laffont and Tirole (1988), to name but a few. However, the majority of these analyze the design of optimal compensation contracts in the presence of moral hazard or adverse selection problems, without addressing the principal's decision whether to retain or replace agents. The latter in turn has been studied in Radner (1986), and Banks and Sundaram

(1998). However, in these models the compensation contract is taken as exogenously given, and the focus lies entirely on the optimal retention rule. In particular, the only incentive mechanism is the threat of being fired, which plays only a marginal role in our paper. To the best of our knowledge, our paper is the first to combine dynamic adverse selection with the theory of learning by experimentation to simultaneously address the questions of optimal contract design and optimal retention. We thus add to the existing literature by showing how the principal can use compensation contracts not only as incentive mechanisms but also to control the informativeness of observed outcomes. Optimal contracts are determined by the trade-off between maximizing the efficiency of the retention rule on one hand, and maximizing short-run expected returns on the other.

We classify compensation schemes according to the risk of the portfolio chosen by the agent. More specifically, by an “aggressive” contract we mean one that induces agents to select high-risk portfolios, while a “conservative” contract induces agents to invest in lower-risk assets. In the context of the compensation packages observed in reality, the more convexity is contained in a contract, the more aggressive it is. Myopically, in the absence of any long-run benefits from learning, the optimal contract will most likely be conservative. In other words, aggressive schemes can be optimal *only* on grounds of their potential informational advantages, *despite* their myopic inferiority.

In our model, the problem of moral hazard is largely eliminated, since the assumed return distribution permits writing contracts that determine the agent’s actions unambiguously. However, the principal faces a problem of *adverse selection* since the agent’s type is ex-ante unknown. We show that “separating” contracts, which are accepted only by good agents and rejected by bad ones, are generally not feasible. In other words, any contract designed to be accepted by good agents will necessarily also be attractive to bad ones. Thus, the first-best

outcome for the principal is generally not implementable. In fact, while in equilibrium good agents will receive no more than their reservation utility, bad agents may extract additional rents. In other words, to ensure that good agents are willing to accept the contract, the principal faces agency costs in form of paying “too much” to bad agents.

As a consequence, good agents act *myopically*, despite the risk of being fired. We believe this to be an accurate albeit slightly stylized description of the observation that investment funds or banks are known to be very reluctant to disclose the reasons for replacing a portfolio manager, since this would cast a bad light on their own reputation. Indeed, Jensen and Murphy (1990) report that “penalties associated with dismissals are very small”. However, since bad agents can extract extra rents from being employed by the principal, the threat of being fired creates an incentive for bad agents to mimic good ones. In order to deter bad agents from doing so, they must be compensated for the loss they incur when fired. This constitutes additional agency costs the principal will have to bear.

Due to the informational asymmetry between good and bad agents, the behaviour of bad agents will be more sensitive with respect to the type of compensation package offered. More specifically, an aggressive contract is most likely to make bad agents take on excessive risk, while the behaviour of good agents is largely unaffected. But bad agents, by definition, are the ones most likely to produce inferior returns when investing aggressively. Therefore, if it was not for any potential informational benefits, the principal would always be better off choosing a conservative contract. On the other hand however, bad agents are more likely to reveal their type when they invest aggressively. In fact we show that the more severe the myopic inferiority of aggressive contracts, the higher their informational advantage. Thus, the principal must trade off short-run utility maximization against long-run informational benefits.

We show that the dynamics of the agent's type and the principal's beliefs can be described by a Markov process on a discrete state space. This enables us to quantify explicitly the informational benefits of different contracts, and to establish necessary and sufficient conditions for certain types of contracts to be optimal. We derive comparative statics results relating the type of optimal contract to various parameters of the model: First, aggressive contracts are *more* likely to be optimal the higher the unconditional expected return on the risky assets, since this makes the aggressive scheme "less costly" in terms of expected utility, while its informational content remains unchanged. Second, aggressive contracts are *less* likely to be optimal the higher the level of the principal's risk-aversion, since risk-aversion increases the utility loss of aggressive contracts. Finally, aggressive contracts are *more* likely to be optimal the higher the likelihood of picking a good agent at random. In other words, if it is easy to find good agents, it is worth bearing the cost of experimentation to detect and fire bad agents as soon as possible.

Although a formal treatment of the dynamic case in which the principal is allowed to change the offered contract is deferred to a subsequent paper, we believe some of the the results to be worth mentioning here. As discussed earlier, aggressive contracts are more useful when the uncertainty about the agent's true type is high. Therefore, the optimal dynamic compensation scheme typically starts off with an aggressive contract and becomes more conservative as the principal grows more confident regarding the agent's type.

The predictions of this paper are consistent with the empirical evidence: First, our results suggest that compensation schemes should be more aggressive in market segments in which returns are more risky. Also, compensation packages should be more aggressive whenever there is higher uncertainty regarding portfolio managers' true ability. This argument may in fact be reversed to estimate the level of uncertainty perceived by market participants from

the types of contracts observed. Finally, our results suggest that compensation schemes should be more aggressive for more junior portfolio managers and become increasingly conservative with increasing seniority. This is consistent with the observation that bonuses make up a much higher fraction of total salary packages for younger portfolio managers, while senior managers often receive very high base salaries and their bonuses are often guaranteed. Moreover, consistent with the predictions of our paper, turnover is much higher amongst junior portfolio managers.

The remainder of this paper is organized as follows; Section 2 outlines the model and introduces the notation. In Section 3, we characterise the dynamics of the principal’s beliefs under different assumptions regarding the agent’s strategies. Conditions under which these types of strategies are indeed optimal are derived in Section 4. Combining the results of the two preceding sections, we characterize optimal first-best and second-best contracts in Section 5, and establish conditions under which experimentation (deviations from the myopic optimum) is optimal. Section 6 concludes.

2 The Model

Trading takes place at discrete times $t = 0, 1, 2, \dots$. By “period t ” we mean the time interval from time t up to time $t + 1$. Two assets are traded, a risk-free “bond” which yields a deterministic gross return R_0 per period, and a risky asset. In any given period t , the risky asset yields a random return $R(t)$ which is realized at the end of the period. The distribution of $R(t)$ depends on the outcome of some exogenous factor $S(t)$. While $S(t)$ might be known to some agents at time t , it is not publically observable and cannot be contracted upon.

Since $S(t)$ conveys information about the distribution of $R(t)$, we will refer to it also as the period- t “signal”. For expositional clarity, we will consider a very simple specification of the return distribution and information structure. While the results developed in this paper can be shown to hold even in much more general scenarios, we believe that the simple specification outlined below allows better insight into the mechanisms behind those results.

2.1 Information Structure and Return Distribution

In any given period t , the gross return $R(t)$ on the risky asset takes on one of three given values, R_- , R_0 , or R_+ , with $R_- < R_0 < R_+$. The signal $S(t)$ on the other hand is assumed to take on either of two values, S_- or S_+ . We call S_- the “bad signal” and S_+ the “good signal”. For $k \in \{-, +\}$ denote by q_k the unconditional probability that $S(t) = S_k$, and for $j \in \{-, 0, +\}$ let q_j^k be the probability that $R(t) = R_j$, conditional on $S(t) = S_k$. To simplify matters further, we make the

Assumption 2.1 *A good outcome can only occur following a good signal, while a bad outcome is only possible when the signal was bad. Formally,*

$$q_-^+ = 0, \quad \text{and} \quad q_+^- = 0.$$

Consequently, we can represent the distribution of returns by a simple two-stage binomial lattice (Figure 1). This assumption in particular implies that observing either $R(t) = R_-$ or $R(t) = R_+$ reveals the signal with certainty. Although this is a somewhat “extreme” feature of the model, preliminary results using “trinomial” branching indicate that the qualitative conclusions of this paper remain true also in more general scenarios.




Figure 1: *Binomial Model of Return Distribution*

To conclude this section, let q_j^* for $j \in \{-, 0, +\}$ denote the *unconditional* probability of outcome $R(t) = R_j$. Evidently,

$$q_+^* = q_+ \cdot q_+^+, \quad q_0^* = q_+ \cdot q_0^+ + q_- \cdot q_0^-, \quad \text{and} \quad q_-^* = q_- \cdot q_-^-.$$

Let $ER(t)$ denote the unconditional expected return in period t . Note that we do not necessarily assume $ER(t) > R_0$. In other words, we do not require a risk-premium on “unhedged” returns. The intuition for this is that the downside risk inherent in $R(t)$ can be avoided by informed agents upon observing $S(t)$.

2.2 Agents

There is a continuum of risk-neutral agents, and any given agent can be either of type θ_b or of type θ_g . We refer to agents of type θ_g as “good” agents and to those of type θ_b as “bad” agents. Denote by π_0 the *unconditional* probability that any randomly chosen agent is of type θ_g . The difference in agents’ types concerns their skill in acquiring information about the return on the risky asset. We think of good agents as having better access to private information or as possessing superior skills in interpreting public information. We model this by making the

Assumption 2.2 *Good agents observe the signal $S(t)$ at the beginning of each period t , while bad agents do not possess any additional information.*

Consequently, good agents base their investment decisions on the conditional return distribution q_j^k , given the observed signal $S(t) = S_k$, while bad agents “know” only the unconditional distribution q_j^* . While this assumption again may seem extreme, relaxing it makes the model far less tractable. Numerical simulations seem to indicate however that the results of this paper remain true under more general specifications, as long as good agents possess some informational advantage in forming return predictions.

Agents may be hired by the principal to manage the principal’s investment portfolio. If employed in any given period t , an agent chooses a portfolio by investing a fraction $\alpha(t)$ of the available funds in the risky asset and the remaining $1 - \alpha(t)$ in the risk-free asset. We restrict $\alpha(t)$ to take values in $[0, 1]$, that is we do not allow short sales. The portfolio return

realized at the end of the period is hence

$$Z(t) = R_0 + \alpha(t) (R(t) - R_0). \quad (1)$$

We assume that the agent's choice $\alpha(t)$ is not verifiable, so that compensation contracts can be written on returns only. This is consistent with the majority of contracts observed in reality, where compensation schemes for portfolio managers typically consist of a flat base salary plus an optional return-dependent bonus. More specifically, a compensation contract is a non-negative, non-decreasing function $w : \mathbb{R} \rightarrow \mathbb{R}$, so that the compensation received by the agent in period t is given by $w(Z(t))$. Let U_0^i denote the value of the outside option for an agent of type θ_i . Since agents are risk-neutral, they will choose a strategy $\alpha(t)$ such as to maximize expected future income,

$$U^i = E \left[\sum_{t=0}^{T-1} \beta^t w(Z(t)) + \beta^T U_0^i \right], \quad (2)$$

where β is the agent's co-efficient of time preference, and T is the number of periods the agent remains employed (we do not rule out the case $T = \infty$). Note that T is ex-ante random, since the principal's decision whether or not to replace the agent depends on realized returns.

2.3 Principal

Denote by $\theta(t)$ the true type of the agent hired by the principal during period t . The principal cannot directly observe $\theta(t)$, but forms beliefs $\pi(t)$ based on the available information. More specifically, $\pi(t)$ is the subjective probability that $\theta(t) = \theta_g$. At the end of period t , the principal observes the portfolio return generated by the agent, $Z(t)$, and updates beliefs accordingly. The principal may then choose to fire and replace the agent.

The principal's preferences are given by a time-separable, non-decreasing, weakly concave von Neumann-Morgenstern utility function $v : \mathbb{R} \rightarrow \mathbb{R}$. In other words, the principal will choose a compensation scheme and a firing rule such as to maximize expected utility

$$V = E \left[\sum_{t=0}^{\infty} \beta^t v(Z(t)) \right], \quad (3)$$

where β is the principal's co-efficient of time-preference. Note that in a slight abuse of notation, we use the same symbol β for both agent and principal, without postulating that they be necessarily equal. Note also that we do not require the utility function v to be strictly concave, that is we do not rule out the case where the principal is risk-neutral.

3 Dynamics of Beliefs

In this section, we will characterize the dynamics of the principal's beliefs, $\pi(t)$, under various assumptions regarding the strategies chosen by agents of either type. More specifically, for $i \in \{b, g\}$ and $k \in \{-, +\}$ let α_k^i denote the portfolio that would be chosen by an agent of type θ_i after having observed signal S_k . Obviously, since only good agents observe the signal, we have to impose the restriction $\alpha_+^b = \alpha_-^b =: \alpha^b$ on bad agents' strategies. Throughout this section, we will take the agent's strategies as given and analyze the resulting dynamics of beliefs. The circumstances under which certain strategies are in fact optimal will be addressed in the next section. We begin by classifying possible strategy combinations.

Definition 3.1 (Classification of Strategies) *A good agent's strategy (α_+^g, α_-^g) is called “normal” if $\alpha_+^g > \alpha_-^g$, and “standard normal” if $\alpha_+^g = 1$ and $\alpha_-^g = 0$. Moreover, a bad agent's strategy α^b is called . . .*

(a) “aggressive mimicking” if $\alpha^b = \alpha_+^g \neq \alpha_-^g$, and

(b) “conservative mimicking” if $\alpha^b = \alpha_-^g \neq \alpha_+^g$.

Finally, we call a bad agent's strategy “separating” if it is not mimicking.

In other words, a good agent's strategy is normal if a larger fraction of the available funds is invested in the risky asset upon receiving a good signal than in the case of a bad signal. The distinction of bad agents' strategies determines the informativeness of observed returns, which we will discuss in greater detail in the following sections.

3.1 Bayesian Updating of Beliefs

For the remainder of this section we will assume that good agents always play the standard normal strategy, while bad agents follow either of the mimicking strategies. More specifically, by the “Aggressive Regime” we mean the case in which the good agent follows a standard normal and the bad agent an aggressive mimicking strategy, and by “Conservative Regime” the analogous case with conservative mimicking. We will discuss later (see Section 4) the circumstances under which these strategies are indeed optimal. Note that in either regime, the only possible return outcomes are the “primitive” returns, i.e. $Z(t) \in \{R_-, R_0, R_+\}$. For $i \in \{b, g\}$ and $j \in \{-, 0, +\}$ denote by p_j^i the probability of observing outcome $Z(t) = R_j$

conditional on the agent being of type θ_i . Tables 1 and 2 list these probabilities for the two regimes.

Proposition 3.2 (Belief Updating) *Suppose the principal's beliefs at the beginning of period t are $\pi(t)$. The Bayesian posterior, conditional on observing outcome $Z(t) = R_j$, is then given by $T_j\pi(t)$, where the updating operator T_j is defined as*

$$T_j\pi := \frac{\pi}{\pi + \kappa_j(1 - \pi)},$$

with $\kappa_j := p_j^b/p_j^g$, and $\kappa_j = +\infty$ whenever $p_j^g = 0$.

PROOF: Bayes' Rule. □

REMARK: The likelihood ratio κ_j determines the direction in which beliefs are updated. More specifically, $\kappa_j > 1$ implies $T_j\pi < \pi$, while $\kappa_j < 1$ implies $T_j\pi > \pi$. Intuitively, $\kappa_j > 1$ corresponds to a return outcome which is more likely to be generated by a bad agent rather than a good agent ($p_j^b > p_j^g$). Hence, observing such an outcome makes it *less likely* that the agent is of the good type, so that beliefs are corrected downwards ($T_j\pi < \pi$). The converse is true for $\kappa_j < 1$. Also note that $\kappa_j = 0$ implies $T_j\pi = 1$, and $\kappa_j = +\infty$ implies $T_j\pi = 0$. These cases correspond to outcomes which reveal the agent's type with certainty. Tables 1 and 2 summarize the updating mechanics for the two regimes.

3.2 Markovian Dynamics of Beliefs and Types

We wish to describe the dynamics of the system in terms of a Markov process. From Tables 1 and 2 we see that the transition probabilities of beliefs depend not only on current beliefs,

j	p_j^g	p_j^b	κ_j	$T_j\pi$
+	$q_+ \cdot q_+^+$	$q_+ \cdot q_+^+$	$= 1$	$= \pi$
0	$q_+ \cdot q_0^+ + q_-$	$q_+ \cdot q_0^+ + q_- \cdot q_0^-$	< 1	$> \pi$
-	0	$q_- \cdot q_-^-$	$+\infty$	$= 0$

Table 1: *Belief Updating, Aggressive Regime*

j	p_j^g	p_j^b	κ_j	$T_j\pi$
+	$q_+ \cdot q_+^+$	0	$= 0$	$= 1$
0	$q_+ \cdot q_0^+ + q_-$	1	> 1	$< \pi$
-	0	0	—	—

Table 2: *Belief Updating, Conservative Regime*

$\pi(t)$, but also on the agent's true type, $\theta(t)$. Since therefore we cannot expect the process of beliefs alone, $\pi(t)$, to be markovian, we have to consider instead the bivariate process of types *and* beliefs, $(\theta(t), \pi(t))$. To characterize the dynamics of this process, we will proceed in two steps: First, we consider the case in which the principal hires an agent once and for all, i.e. $\theta(t) \equiv \theta(0)$ for all t . Second, we determine the principal's optimal policy of firing and replacing agents and analyse the manner in which this affects the dynamics of beliefs and types.

From Tables 1 and 2 we see that any outcome *other* than $Z(t) = R_0$ either produces no information at all ($T_j\pi = \pi$), or reveals the agent's type with certainty ($T_j\pi \in \{0, 1\}$). Hence, starting from any arbitrary belief π , the only possible values for the updated beliefs are π itself, $T_0\pi$ or the extreme values zero and one. Motivated by this observation, we define

$$\pi_n := T_0^n \pi_0 = \frac{\pi_0}{\pi_0 + (\kappa_0)^n (1 - \pi_0)} \quad \text{for } n = 0, 1, \dots$$

A suitable state space on which beliefs “live” is hence given by $E := \{0, 1\} \cup \{\pi_0, \pi_1, \dots\}$. Note that the π_n will be different in the two regimes. More precisely, we find that $\pi_{n+1} > \pi_n$ in the aggressive regime, while $\pi_{n+1} < \pi_n$ in the conservative regime. The intuition behind this is straight-forward: In the conservative regime, outcome $Z(t) = R_0$ is more likely to come from a bad agent, hence beliefs are shifted towards the bad type. Conversely, in the aggressive regime, outcome $Z(t) = R_0$ is more likely to be produced by a good agent, and beliefs are accordingly adjusted upwards.

We show in the appendix (Proposition A.1) that the process $(\theta(t), \pi(t))$ of types and beliefs is indeed a Markov process on the state space $\{\theta_b, \theta_g\} \times E$. The transition probabilities are depicted in Figures 2 and 3. Evidently, the most informative event in the conservative regime is observing a high return $Z(t) = R_+$. This is so because in this regime bad agents invest conservatively and will hence always produce a portfolio return of R_0 . Thus, observing a high return reveals with certainty that the agent is of the good type. Conversely, in the aggressive regime, observing a low return $Z(t) = R_-$ is most informative because a good agent would never have invested in the risky asset upon receiving a bad signal. In other words, in the conservative regime it is the *good* agents who reveal their type by producing high returns, while in the aggressive regime it is the bad agents revealing their type through poor performance.

It should be noted that although revelation with *certainty* is a consequence of the discrete return distribution, the qualitative properties of different regimes will remain the same even under a continuous distribution: it will always be the case that high returns in a conservative regime indicate a good agent, while low returns in an aggressive regime indicate a bad one. However, we view the existence of separating regimes as a pure artifact of the discrete distribution which is not only extremely unrealistic but also not robust with respect to

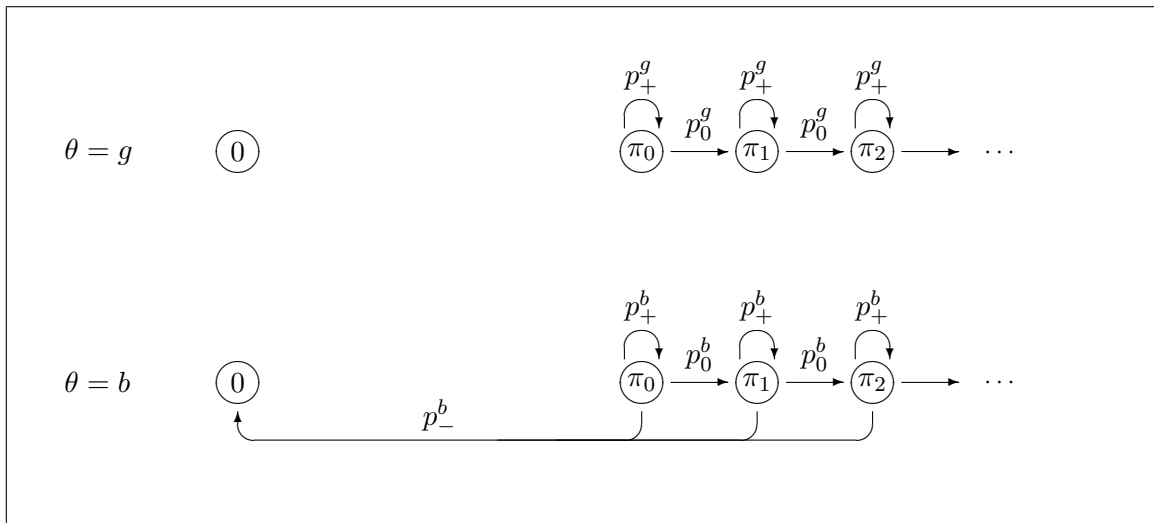


Figure 2: *Dynamics of Beliefs and Types, Aggressive Regime*

different distributional assumptions. Therefore, we discard the separating regimes and focus entirely on the mimicking regimes.

3.3 Hiring and Firing

In either regime, returns generated by good agents are always at least as high as those generated by bad agents. Therefore, we expect the principal's value function to be increasing in current beliefs $\pi(t)$, (for a formal proof, see Proposition 5.4.) On the other hand, whenever a new agent is hired, beliefs will be reset to the initial prior, π_0 . Therefore, it is optimal for the principal to fire and replace the agent if and only if the posterior belief after observing the most recent return falls below π_0 .

We show in the appendix (Proposition A.2) that if the principal follows this firing strategy, the process $(\theta(t), \pi(t))$ of types and beliefs is indeed a Markov process. The transition

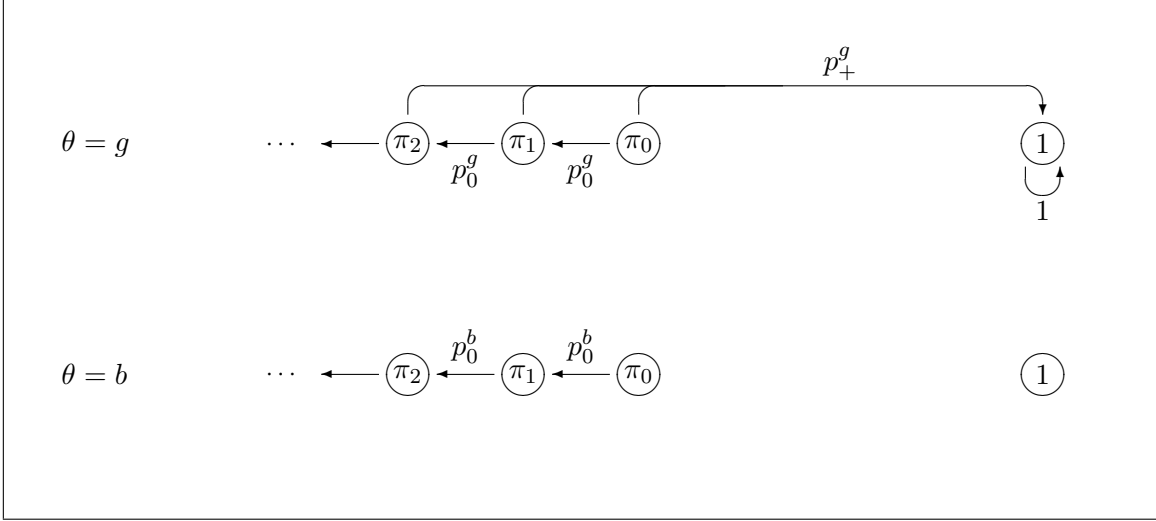


Figure 3: *Dynamics of Beliefs and Types, Conservative Regime*

probabilities are depicted in Figures 4 and 5.

We are now ready to state the main result of this section, which characterizes the dynamics of *expected* beliefs. Throughout the remainder of the paper, we will use an index $m \in \{a, c\}$ to indicate either the aggressive ($m = a$) or the conservative ($m = c$) regime, respectively.

Theorem 3.3 *In regime m , the unconditional expectation of the principal's beliefs can be expressed in the form $E[\pi(t)] = 1 - \lambda_m^t (1 - \pi_0)$, where*

$$\begin{aligned} \lambda_a &= 1 - q_-^* \pi_0 \quad \text{in the aggressive regime, and} \\ \lambda_c &= 1 - q_+^* \pi_0 \quad \text{in the conservative regime.} \end{aligned}$$

PROOF: See Appendix A.1. □




Figure 4: *Dynamics of Beliefs and Types, Aggressive Regime with Firing/Hiring*

Corollary 3.4 *Since $\pi(t)$ is the conditional distribution of $\theta(t)$ given the principal's information, the law of iterated expectations implies $P[\theta(t) = \theta_g] = 1 - \lambda_m^t (1 - \pi_0)$.*

The co-efficient λ_m measures the benefit from learning: the lower λ_m , the shorter the average time it takes the principal to find a good agent. From the definition of λ_m , we can identify two distinct effects: First, λ_m is lower the higher the probability of hiring a good agent, π_0 . This simply reflects the fact that it is easier to pick a good agent by chance the higher the fraction of good agents is in the total population. Second, λ_m is also lower the higher the probability of the event that reveals the agent's type with certainty, (q_-^* in the aggressive regime, and q_+^* in the conservative regime.) In other words, the aggressive regime performs better in finding a good agent whenever q_-^* is large in comparison with q_+^* , that is whenever low returns are more likely than high returns. On the other hand, the higher the probability of low returns, the higher the principal's loss of expected per-period utility when bad agents invest aggressively. Thus, the principal must trade off the informational advantage of the aggressive

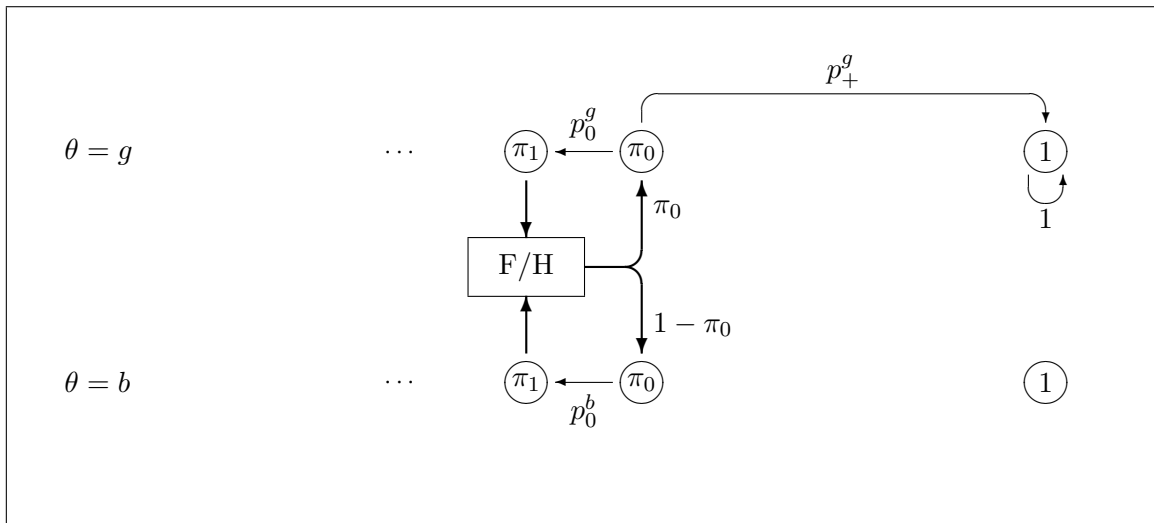


Figure 5: *Dynamics of Beliefs and Types, Conservative Regime with Firing/Hiring*

regime against its inferiority in terms of myopic utility maximization. This fundamental trade-off, typical for models of learning by experimentation, will determine which type of regime is optimal.

4 Compensation Contracts and Portfolio Selection

In this section, we characterize the optimal strategies for agents of either type. We will establish necessary and sufficient conditions regarding the compensation scheme offered by the principal which ensure that certain strategy combinations are optimal. From the preceding section, we know that for a fixed compensation scheme, the probability of being fired does not depend on the principal's beliefs. Therefore, we can restrict our analysis to *stationary* strategies, i.e. those which depend only on the agent's type and the observed signal. In other words, a good agent's strategy is entirely characterised by the pair (α_+^g, α_-^g) of port-

folio choices conditional on the observed signal, while a bad agent's strategy is simply given by the unconditional portfolio choice α^b . For what follows, we fix a given compensation scheme w . For $k \in \{-, +\}$, denote by $\bar{w}_k(\alpha)$ the expected compensation if the agent chooses portfolio α , conditional on signal S_k . Evidently,

$$\bar{w}_k(\alpha) = q_k^k \cdot w(R_0 + \alpha(R_k - R_0)) + q_0^k \cdot w(R_0).$$

Finally, denote by $\bar{w}(\alpha)$ the unconditional expected compensation,

$$\bar{w}(\alpha) = q_+ \cdot \bar{w}_+(\alpha) + q_- \cdot \bar{w}_-(\alpha).$$

4.1 Good Agents

Let us fix a given strategy α^b followed by the bad agent. We wish to characterize the optimal strategy for the good agent, given the compensation scheme offered by the principal. Denote by $U^g(\alpha_+, \alpha_-)$ the good agent's value function for a given strategy (α_+, α_-) , and let U_0^g denote the value of the good agent's "outside option".

Proposition 4.1 *The good agent's value function solves the functional equation*

$$U^g(\alpha_+, \alpha_-) = q_+ \cdot \bar{w}_+(\alpha_+) + q_- \cdot \bar{w}_-(\alpha_-) + \beta \{ \nu U_0^g + (1 - \nu) U^g(\alpha_+, \alpha_-) \}, \quad (4)$$

where ν denotes the probability of being fired in any given period.

REMARK: The firing probability ν depends of course also on the strategies chosen by both types of agent. More specifically, $\nu = p_0^g$ whenever $\alpha_+ = \alpha^b$, and $\nu = 0$ otherwise.

Obviously, the agent will only be willing to work for the principal if the expected utility of doing so is at least as high as the agent’s outside option, U_0^g . In other words, the compensation scheme offered by the principal must satisfy the *participation constraint*

$$\sup_{\alpha_+, \alpha_-} U^g(\alpha_+, \alpha_-) \geq U_0^g. \quad (5)$$

On the other hand, the principal will seek to find the “cheapest” contract which satisfies this constraint, so that the constraint will be binding at the optimum. In other words, in equilibrium good agents will be indifferent between the outside option and working for the principal. We are now ready to characterize the good agent’s optimal strategy.

Theorem 4.2 *Suppose the compensation scheme is chosen such that*

$$q_+ \sup_{\alpha_+} \bar{w}_+(\alpha_+) + q_- \sup_{\alpha_-} \bar{w}_-(\alpha_-) \equiv (1 - \beta)U_0^g. \quad (6)$$

Then the participation constraint (5) is satisfied with equality, and any α_+^ and α_-^* that attain the suprema in the above expression are optimal, i.e.*

$$U^g(\alpha_+^*, \alpha_-^*) = \sup_{\alpha_+, \alpha_-} U^g(\alpha_+, \alpha_-) = U_0^g.$$

PROOF: See Appendix A.2. □

This theorem simply states that good agents will act *myopically*, i.e. choose a strategy (α_+, α_-) such as to maximize the expected per-period compensation. In other words, the threat of being fired is irrelevant for choosing an optimal strategy. While this is not critical for our analysis, we believe it to be an accurate albeit somewhat stylized description of the real world, where banks or investment funds are known to be extremely reluctant to disclose the reasons for replacing a portfolio manager. Indeed, Jensen and Murphy (1990) report that “penalties associated with dismissals are very small”.

4.2 Bad Agents

Analogous to the preceding section, we fix an arbitrary strategy (α_+, α_-) for the good agent, and denote by $U^b(\alpha)$ the bad agent's value function, and by U_0^b the value of a bad agent's outside option. The first result is similar to the one obtained for good agents:

Proposition 4.3 *The bad agent's value function satisfies the functional equation*

$$U^b(\alpha) = \bar{w}(\alpha) + \beta \{ \nu U_0^b + (1 - \nu) U^b(\alpha) \},$$

where ν is the probability of being fired.

REMARK: As for the good agent, the firing probability ν depends on the strategies chosen by both types of agent. More specifically, $\nu = p_-^b$ whenever $\alpha = \alpha_+^g$, and $\nu = 1$ otherwise.

Obviously, a given contract will be accepted by bad agents if and only if it satisfies the corresponding *participation constraint*,

$$\sup_{\alpha} U^b(\alpha) \geq U_0^b. \tag{7}$$

However, while any contract must necessarily satisfy the participation constraint for good agents, this is not the case for bad agents. In fact, the principal would prefer to write a contract which is accepted only by good agents but rejected by bad ones. However, it is a priori not clear whether such contracts always exist. In fact we will show that if the values of the agents' outside options are sufficiently different across types, then the only feasible contracts are those which are accepted by both types. Moreover, in this case only the participation constraint for good agents will be binding. Conversely, if the participation constraint for bad agents is binding, the contract would not be accepted by good agents.

In other words, the principal incurs agency costs due to the problem of adverse selection. Before addressing this issue in more detail in the next section (Theorem 5.2), we need to characterize bad agents' optimal strategies:

Theorem 4.4 *Suppose the compensation scheme is chosen such that*

$$\sup_{\alpha} \bar{w}(\alpha) \geq (1 - \beta)U_0^b. \quad (8)$$

Then the participation constraint (7) for bad agents is satisfied, and any α^ that attains the supremum is optimal if and only if*

$$\bar{w}(\alpha^*) \geq \bar{w}(\alpha_+^g) + \beta(1 - p_-^b) \{ \bar{w}(\alpha^*) - (1 - \beta)U_0^b \}. \quad (9)$$

Otherwise, α_+^g is optimal.

PROOF: See Appendix A.3. □

Note that the last term in (9) can be interpreted as the loss the bad agent incurs by choosing *not* to mimic the good agent and thus reducing the chance of remaining employed. In other words, to deter bad agents from mimicking good ones, they have to be compensated for the utility loss associated with the threat of being fired. This constitutes additional agency costs that have to be borne by the principal.

We have seen that agents' optimal strategies are entirely determined by the conditional expected compensations $\bar{w}_k(\alpha)$. Conversely, it is easy to see that for any arbitrary functions $f_k(\alpha)$, a contract w can be constructed such that $\bar{w}_k(\alpha) \equiv f_k(\alpha)$. Figures 6 and 7 show examples of contracts that implement the aggressive respectively conservative regime. Note however that the depicted contracts need not represent the best possible way to implement

the respective regime. In fact, the principal will choose contracts that implement the optimal regime while minimizing the expected cost. To relate our results to the types of compensation schemes observed in reality, we broadly classify these schemes into

affine	$w(z) = a_0 + a_1z$	can be either (typically conservative)
simple bonus	$w(z) = a_0 + b(z - z_0)^+$	always aggressive
mixed	$w(z) = a_0 + a_1z + b(z - z_0)^+$	can be either (typically aggressive)

In particular, affine contracts are aggressive if and only if $ER(t) > R_0$, and conservative otherwise. Mixed contracts on the other hand can only be conservative if $ER(t) < R_0$. In other words, a substantial fraction of the bonus-based compensation schemes observed in reality classify as aggressive in the context of our model.

5 Optimal Contracts

In this section, we will study the circumstances under which certain types of contracts are optimal from the principal's vantage point. As discussed earlier, we believe separating regimes to be an artifact of the discrete return distribution and as such neither realistic nor robust with respect to changes in the distribution. Therefore, we will discard these regimes and focus entirely on the two mimicking regimes. As before, we use an index $m \in \{a, c\}$ to indicate either the aggressive ($m = a$) or the conservative ($m = c$) regime, respectively. For $i \in \{b, g\}$ and $m \in \{a, c\}$, let \bar{v}_m^i denote the principal's expected per-period utility in regime m , given that the agent is of type θ_i . Evidently,

$$\begin{aligned} \bar{v}_m^g &= q_+^*v(R_+) + (1 - q_+^*)v(R_0) =: \bar{v}^g \quad \text{for both } m = a \text{ and } m = c, \text{ while} \\ \bar{v}_a^b &= q_+^*v(R_+) + q_0^*v(R_0) + q_-^*v(R_-), \quad \text{and} \quad \bar{v}_c^b = v(R_0). \end{aligned}$$

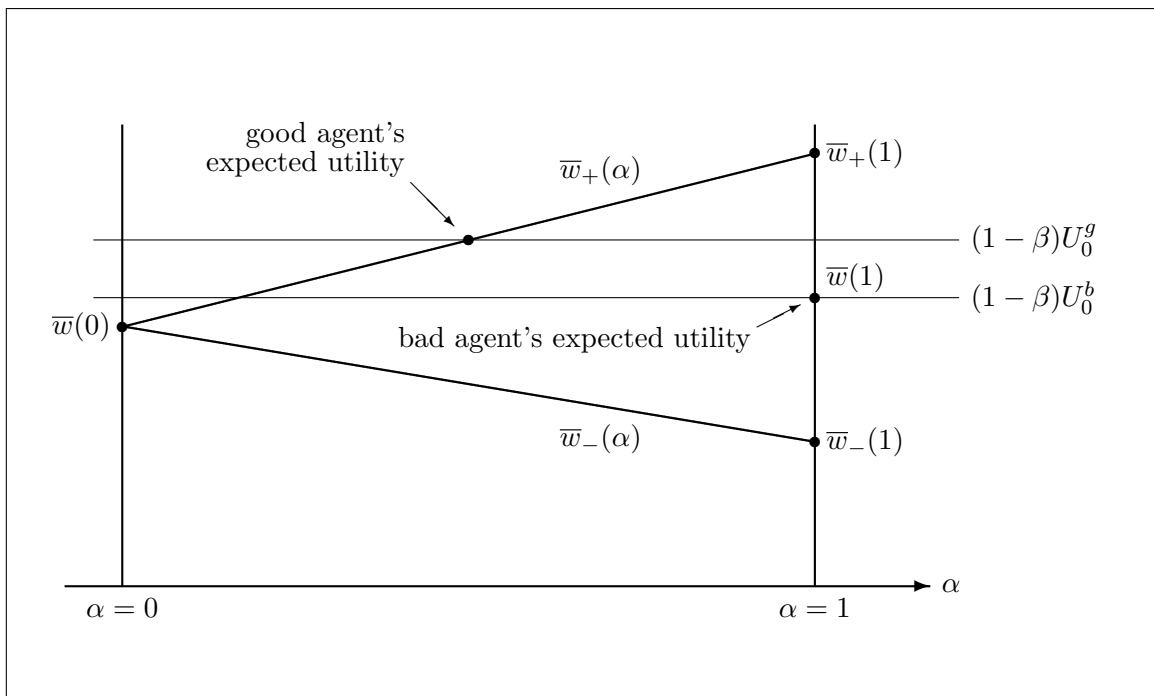


Figure 6: *Example of an Aggressive Contract*

Note that it will always be the case that $\bar{v}^g > \bar{v}_m^b$ for either regime m . Moreover, if the principal is sufficiently risk-averse, or if $ER(t) < R_0$, we will also have $\bar{v}_a^b < \bar{v}_c^b$. This is in fact the most interesting case, since it implies that myopically, conservative contracts are always optimal. In other words, in this case aggressive contracts may be optimal on grounds of their informational benefits only, *despite* their myopic inferiority.

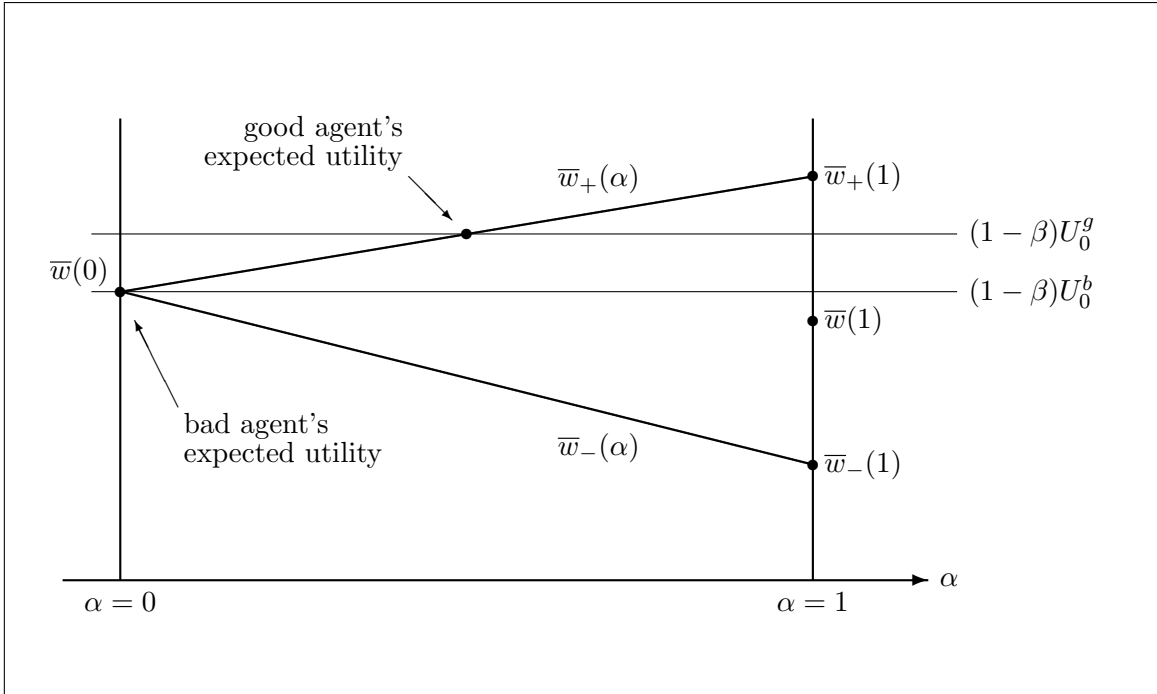


Figure 7: *Example of a Conservative Contract*

5.1 First-Best

Before characterizing the optimal second-best contract, it is helpful to analyse the principal's value function in the first-best scenario.

Proposition 5.1 *Suppose the principal can observe the agent's type ex-ante. In this case, only good agents are hired, and the principal's value function is given by*

$$V^* = \frac{\bar{v}^g}{1 - \beta}. \quad (10)$$

We will investigate below the circumstances under which it is possible to implement the first-best outcome even if the agent's type is not observable. More specifically, we will see that in some cases contracts may be written that are accepted by good agents but rejected by bad ones. In general however, we cannot expect such contracts to be always feasible. More specifically,

Theorem 5.2 *The first-best outcome can be implemented if and only if*

$$U_0^g < (1 + q_-)U_0^b. \tag{11}$$

PROOF: See Appendix A.4. □

In other words, if the value of agents' outside option is sufficiently different across types, only second-best contracts are feasible. More specifically, if (11) is violated, any non-negative compensation scheme which satisfies the participation constraint for good agents, (5), necessarily also satisfies that for bad agents, (7). In fact, in equilibrium the participation constraint will be binding only for good agents, so that bad agents can extract additional rents at the expense of the principal. The key assumption here is non-negativity: If the principal was allowed to write contracts that specify negative compensation at least in some cases, the first-best outcome could always be implemented. However, since compensation packages that require employees to pay the employer are rarely observed in reality, we maintain the non-negativity restriction. It is therefore essential to ask what types of second-best contracts are optimal in cases where the first-best contract is not feasible.

5.2 Second-Best

We begin our analysis of the second-best case by characterizing the principal's value function in the absence of learning. In this case, the principal faces the agency cost associated with the unobservability of the agent's type, without making use of the information contained in observed returns to lower these cost.

Proposition 5.3 *Without learning, the principal's total expected utility in regime m is*

$$V_m^0 = V^* - A_m^0 \quad \text{with} \quad A_m^0 = \frac{1 - \pi_0}{1 - \beta} \{ \bar{v}^g - \bar{v}_m^b \}. \quad (12)$$

In particular, the optimal contract in this case is aggressive if and only if $\bar{v}_a^b \geq \bar{v}_c^b$, and conservative otherwise.

Thus, the principal's value function in the second-best case can be decomposed into the first-best value function, V^* , minus the term A_m^0 , which we may interpret as *agency cost*. Note first that A_m^0 is always positive in either regime, so that the first-best outcome indeed dominates the second-best. Intuitively, learning should *reduce* agency cost; indeed we find

Proposition 5.4 *With learning, the principal's total expected utility in regime m is*

$$V_m = V^* - A_m \quad \text{with} \quad A_m = \frac{1 - \pi_0}{1 - \beta \lambda_m} \{ \bar{v}^g - \bar{v}_m^b \}, \quad (13)$$

where $\lambda_a = 1 - q_-^ \pi_0$ and $\lambda_c = 1 - q_+^* \pi_0$ are defined as in Theorem 3.3.*

PROOF: See Appendix A.5. □

Note that the agency cost in this case, A_m , differ from the agency cost in the case without learning, A_m^0 , only by the co-efficient λ_m . From Section 3 we know that λ_m measures the benefit from learning: the higher the speed of learning, the lower λ_m , and the greater the reduction of agency cost. In particular, $\lambda_m = 1$ corresponds to no learning at all. The size of the agency cost is hence determined by trading off the myopic utility cost, $\bar{v}^g - \bar{v}_m^b$, against the benefits from learning measured by λ_m .

A given regime m is said to be *myopically optimal* if it minimizes A_m^0 , and *optimal* if it minimizes A_m . Since the speed of learning varies across regimes, it is possible that the overall optimal regime is different from the myopically optimal one. In other words, it may be optimal for the principal to deviate from the myopic optimum whenever the cost of doing so is outweighed by the long-run informational benefits. Such cases are referred to as *optimal experimentation*. Obviously, the optimal second-best contract will be aggressive if and only if $V_a > V_c$, and conservative otherwise. Using Proposition 5.4, we find

Theorem 5.5 *The optimal second-best contract is aggressive if and only if*

$$\frac{1 - \beta\lambda_a}{1 - \beta\lambda_c} > \frac{\bar{v}^g - \bar{v}_a^b}{\bar{v}^g - \bar{v}_c^b}, \quad (14)$$

and conservative otherwise. The principal's value function is given as in Proposition 5.4.

Note first that the right-hand side of (14) is greater than one if and only if the conservative regime is myopically optimal ($\bar{v}_c^b > \bar{v}_a^b$), and smaller than one otherwise. Conversely, the left-hand side of (14) is greater than one if and only if aggressive contracts possess an informational advantage over conservative ones ($\lambda_a < \lambda_c$), and smaller than one otherwise. The intuitive interpretation of condition (14) is straight-forward: Suppose for example that the conservative contract is myopically optimal. In this case the right-hand side of (14) measures

the relative utility loss incurred by deviating from the myopic optimum, while the left-hand side measures the informational benefits from doing so. In other words, aggressive contracts are optimal if and only if the informational gain outweighs the cost of experimentation.

Obviously, as the principal's level of risk-aversion increases, the right-hand side of (14) becomes larger, so that aggressive contracts are *less likely* to be optimal. Intuitively, risk-aversion increases the utility loss incurred when bad agents invest aggressively, while leaving the informational aspects unchanged. Moreover, an increase in the proportion π_0 of good agents in the total population will widen the gap between the co-efficients of informativeness, λ_m . In particular, if aggressive contracts do possess an informational advantage, an increase in π_0 would amplify this advantage, thus making aggressive contracts *more likely* to be optimal. On the other hand, the effects of changes in the distribution of returns are not quite so clear cut. For example, a shift in probability from q_-^* to q_+^* , (increasing the expected return on the risky asset,) will causes both sides of (14) to decrease. Whether this change goes in favour of aggressive or conservative contracts depends on the other parameters of the model. However, increasing the expected return on the risky asset by raising the primitive returns R_j , has no effect on the informativeness of different regimes while reducing the cost of experimentation, thus making the aggressive regime *more likely* to be optimal.

6 Conclusion and Outlook

We show how the bonus-based compensation schemes for portfolio managers observed in reality can be justified theoretically, despite the fact that they are known to induce excessive risk-taking. More specifically, we characterize such bonus schemes endogenously as opti-

mal contracts in an infinite-horizon principal-agent model in which the principal may learn about the agent's ex-ante uncertain type by means of "optimal experimentation". The type of optimal contract is determined by the fundamental trade-off between short-run utility maximization and long-run informational benefits, which is typical for models in the theory of learning by experimentation.

The results we derive are consistent with many of the empirical findings in the area of fund management compensation. In particular, we provide an alternative theoretical explanation for the surprising amount of convexity contained in most compensation packages observed in reality that is not based on agents' risk-aversion or dis-utility of effort. Our explanation is consistent with the observation that management turnover is inversely related to performance, and that the variance of returns is increased in the period prior to the dismissal.

As an obvious and important extension to the work presented in this paper, we address in a forthcoming paper the case in which the principal is allowed to change the compensation package offered to managers according to the evolution of the principal's beliefs. Preliminary results indicate that in addition to the static regimes considered here, there will be a "switching regime" in which the optimal contract is aggressive at first and becomes increasingly conservative as the principal becomes sufficiently confident regarding the agent's type. These results explain why compensation contracts are increasingly conservative as fund managers become more senior, and why management turnover is higher amongst more junior managers. Furthermore, we will study the case in which the principal can decide what fraction of the available funds to allocate to the agent. Intuitively, the possibility of reducing the stakes during the experimentation phase should reduce the cost of experimentation, thus making aggressive contracts even more attractive.

A Mathematical Appendix

On some suitably chosen probability space, let $(S(t), R(t))_{t=0,1,\dots}$ be a sequence of serially independent bivariate random variables, taking values in $\{S_-, S_+\} \times \{R_-, R_0, R_+\}$, with

$$\text{Prob}(S(t) = S_k; R(t) = R_j) = q_k \cdot q_j^k; \quad \text{for } k \in \{-, +\} \text{ and } j \in \{-, 0, +\};$$

and let $(\vartheta(t))_{t=0,1,\dots}$ be a sequence of serially independent random variables, taking values in $\{\theta_b, \theta_g\}$, with $\text{Prob}(\vartheta(t) = \theta_g) = \pi_0$.

Proposition A.1 *Let $S(t)$, $R(t)$, and $\vartheta(t)$ be defined as above. Set $\pi(0) = \pi_0$ and $\theta(0) = \vartheta(0)$. For $\theta(t)$ and $\pi(t)$ given, we define recursively:*

Let $\alpha(t) = \alpha_k^i$, where $i \in \{b, g\}$ and $k \in \{-, +\}$ are such that $\theta(t) = \theta_i$ and $S(t) = S_k$. Define $Z(t) = R_0 + \alpha(t)(R(t) - R_0)$ and choose $j \in \{-, 0, +\}$ such that $Z(t) = R_j$. Let

$$\begin{aligned} \pi(t+1) &= T_j \pi(t), \\ \theta(t+1) &= \theta(t). \end{aligned}$$

Then $(\theta(t), \pi(t))_{t=0,1,\dots}$ is a Markov process on the state space $\{\theta_b, \theta_g\} \times E$. The transition probabilities are as depicted in Figure 2 for the aggressive regime, and Figure 3 for the conservative regime.

PROOF: Trivial. □

Proposition A.2 *Let $S(t)$, $R(t)$, and $\vartheta(t)$ be defined as above. Set $\pi(0) = \pi_0$ and $\theta(0) = \vartheta(0)$. For $\theta(t)$ and $\pi(t)$ given, we define recursively:*

Let $\alpha(t) = \alpha_k^i$, where $i \in \{b, g\}$ and $k \in \{-, +\}$ are such that $\theta(t) = \theta_i$ and $S(t) = S_k$. Define $Z(t) = R_0 + \alpha(t)(R(t) - R_0)$ and choose $j \in \{-, 0, +\}$ such that $Z(t) = R_j$. Let

$$\begin{aligned}\pi(t+1) &= \begin{cases} \pi_0 & \text{if } T_j\pi(t) < \pi_0, \text{ and} \\ T_j\pi(t) & \text{otherwise.} \end{cases} \\ \theta(t+1) &= \begin{cases} \vartheta(t+1) & \text{if } T_j\pi(t) < \pi_0, \text{ and} \\ \theta(t) & \text{otherwise.} \end{cases}\end{aligned}$$

Then $(\theta(t), \pi(t))_{t=0,1,\dots}$ is a Markov process on the state space $\{\theta_b, \theta_g\} \times E$. The transition probabilities are as depicted in Figure 4 for the aggressive regime, and Figure 5 for the conservative regime.

PROOF: Trivial. □

A.1 Proof of Theorem 3.3

We consider the two regimes separately, omitting the index $m = a, c$.

AGGRESSIVE REGIME: Let $\theta(t)$ be defined as in Proposition A.2. From Figure 4 it is evident that in the aggressive regime, $\theta(t)$ is itself Markov on $\{\theta_b, \theta_g\}$, with transition matrix

$$\begin{pmatrix} \lambda & 1 - \lambda \\ 0 & 1 \end{pmatrix},$$

where $\lambda = 1 - \pi_0 p_-^b$. Let Q be the transpose of the transition matrix. It is easy to see that

$$Q^t = \begin{pmatrix} \lambda^t & 0 \\ 1 - \lambda^t & 1 \end{pmatrix}.$$

Let $\mu(t)$ denote the distribution of $\theta(t)$ on $\{\theta_b, \theta_g\}$. Obviously, $\mu(0) = (1 - \pi_0, \pi_0)'$, and $\mu(t) = Q^t \mu(0) = (\lambda^t(1 - \pi_0), 1 - \lambda^t(1 - \pi_0))'$. Since by definition, $\pi(t)$ is the conditional distribution of $\theta(t)$ given the principal's information, the law of iterated expectations gives

$$E[\pi(t)] = P[\theta(t) = \theta_g] = 1 - \lambda^t(1 - \pi_0),$$

which was to be shown. □

CONSERVATIVE REGIME: Let $\pi(t)$ defined as in Proposition A.2. From Figure 5 it is evident that in the conservative regime, $\pi(t)$ is itself Markov on $\{\pi_0, 1\}$, with transition matrix

$$\begin{pmatrix} \lambda & 1 - \lambda \\ 0 & 1 \end{pmatrix},$$

where $\lambda = \pi_0 p_0^g + (1 - \pi_0) p_0^b$. Note that in the conservative regime, $p_0^b \equiv 1$, so that we can simplify $\lambda = 1 - \pi_0(1 - p_0^g) = 1 - \pi_0 p_+^g$. Let Q denote the transpose of the transition matrix. Analogous to the aggressive regime, we find

$$Q^t = \begin{pmatrix} \lambda^t & 0 \\ 1 - \lambda^t & 1 \end{pmatrix}.$$

Denote by $\mu(t)$ the distribution of $\pi(t)$ on $\{\pi_0, 1\}$. Obviously, $\mu(0) = (1, 0)'$, and $\mu(t) = Q^t \mu(0) = (\lambda^t, 1 - \lambda^t)'$. Consequently,

$$E[\pi(t)] = \lambda^t \cdot \pi_0 + (1 - \lambda^t) \cdot 1 = 1 - \lambda^t \cdot (1 - \pi_0),$$

which was to be shown. □

A.2 Proof of Theorem 4.2

Fix any arbitrary α_+ and α_- . Using Proposition 4.1 we find

$$U^g(\alpha_+, \alpha_-) = \underbrace{q_+ \bar{w}_+(\alpha_+) + q_- \bar{w}_-(\alpha_-)}_{\leq (1-\beta)U_0^g} + \beta \{ \nu U_0^g + (1-\nu)U^g(\alpha_+, \alpha_-) \}.$$

Hence, $(1 - \beta(1 - \nu)) \cdot (U^g(\alpha_+, \alpha_-) - U_0^g) \leq 0$, with equality if and only if

$$q_+ \bar{w}_+(\alpha_+) + q_- \bar{w}_-(\alpha_-) = (1 - \beta)U_0^g.$$

This implies the desired result since $\beta(1 - \nu) < 1$. □

A.3 Proof of Theorem 4.4

Using Proposition 4.3, we find

$$U^b(\alpha^*) = \underbrace{\bar{w}(\alpha^*)}_{\geq (1-\beta)U_0^b} + \beta \{ \nu U_0^b + (1-\nu)U^b(\alpha^*) \}.$$

Hence, $(1 - \beta(1 - \nu)) \cdot (U^b(\alpha^*) - U_0^b) \geq 0$, which implies $U^b(\alpha^*) \geq U_0^b$ since $\beta(1 - \nu) < 1$. To prove the optimality condition, observe first that ν is the same for all $\alpha \neq \alpha_+^g$. It is hence sufficient to compare $U^b(\alpha)$ with $U^b(\alpha_+^g)$ for $\alpha \neq \alpha_+^g$. Using Proposition 4.3 again, we find

$$(1 - \beta(1 - p_-^b)) \cdot (U^b(\alpha_+^g) - U^b(\alpha)) = \bar{w}(\alpha_+^g) + \beta p_-^b U_0^b + (1 - \beta(1 - p_-^b)) \{ \bar{w}(\alpha) - \beta U_0^b \}.$$

In other words, $U^b(\alpha_+^g) \geq U^b(\alpha)$ if and only if

$$\bar{w}(\alpha_+^g) \geq (1 - \beta(1 - p_-^b))\bar{w}(\alpha) + \beta(1 - p_-^b)(1 - \beta)U_0^b.$$

Note that the right-hand side of the above inequality is just a convex combination of $\bar{w}(\alpha)$ and $(1 - \beta)U_0^b$. Hence, if $\alpha^* = \alpha_+^g$, the above inequality is clearly satisfied and α_+^g is optimal. Otherwise, for $\alpha = \alpha^* \neq \alpha_+^g$ the above inequality is the contraposition of (9). □

A.4 Proof of Theorem 5.2

We must have $\bar{w}_-(1) < w(0) < \bar{w}_+(1)$ if good agents are to follow the standard normal strategy. Due to the non-negativity assumption, we may without loss of generality assume $\bar{w}_-(1) = 0$. We will show that if $U_0^g \geq (1 + q_+)U_0^b$, then the participation constraint for good agents, (5), implies that for bad agents, (7). We hence assume

$$q_+\bar{w}_+(1) + (1 - q_+)w(0) \geq (1 - \beta)U_0^g.$$

From Theorem 4.4 we know that (7) is satisfied at $\alpha = 0$ when $w(0) \geq (1 - \beta)U_0^b$. Conversely, if $w(0) < (1 - \beta)U_0^b$, we find

$$\begin{aligned} q_+\bar{w}_+(1) &\geq (1 - \beta)U_0^g - (1 - q_+)w(0) \\ &\geq (1 + q_+)(1 - \beta)U_0^b - (1 - q_+)(1 - \beta)U_0^b = (1 - \beta)U_0^b, \end{aligned}$$

so that the participation constraint for bad agents is satisfied at $\alpha = 1$.

Conversely, it is easy to see that if $U_0^g < (1 + q_+)U_0^b$, a non-negative contract can be constructed that satisfies the participation constraint for good agents, (5), but violates that for bad agents, (7). This completes the proof. \square

A.5 Proof of Proposition 5.4

We consider the two regimes separately, omitting the index $m = a, c$.

AGGRESSIVE REGIME: From the proof of Theorem 3.3 we know that in the aggressive regime, the system is entirely described by the agent's current type, $\theta(t)$. Denote by $V(\theta_g)$

respectively $V(\theta_b)$ the principal's value function in either state. Using the Markov transition matrix derived in the proof of Theorem 3.3 we find

$$\begin{aligned} V(\theta_g) &= \bar{v}^g + \beta V(\theta_g) \\ V(\theta_b) &= \bar{v}^b + \beta \{ \lambda V(\theta_b) + (1 - \lambda) V(\theta_g) \} \end{aligned}$$

Combining these two equations it is easy to see that

$$V(\theta_b) - V(\theta_g) = \frac{\bar{v}^b - \bar{v}^g}{1 - \beta\lambda}.$$

Since initially, the distribution of the agent's type is given by π_0 , the principal's overall value function is hence given by

$$V(\theta_g) + (1 - \pi_0) (V(\theta_b) - V(\theta_g)) = \frac{\bar{v}^g}{1 - \beta} + \frac{1 - \pi_0}{1 - \beta\lambda} (\bar{v}^b - \bar{v}^g),$$

which was to be shown. □

CONSERVATIVE REGIME: From the proof of Theorem 3.3 we know that in the conservative regime, the system is entirely described by the principal's belief, $\pi(t)$, which can only take on the values π_0 or one. Denote by $V(\pi_0)$ respectively $V(1)$ the principal's value function in either state. Analogous to the aggressive case we find

$$\begin{aligned} V(1) &= \bar{v}^g + \beta V(1) \\ V(\pi_0) &= \pi_0 \bar{v}^g + (1 - \pi_0) \bar{v}^b + \beta \{ \lambda V(\pi_0) + (1 - \lambda) V(1) \} \end{aligned}$$

Combining these two equations it is easy to see that

$$V(\pi_0) - V(1) = \frac{(1 - \pi_0)(\bar{v}^b - \bar{v}^g)}{1 - \beta\lambda}.$$

Since initially, the distribution of the agent's type is given by π_0 , the principal's overall value function is simply given by

$$V(\pi_0) = V(1) + (V(\pi_0) - V(1)) = \frac{\bar{v}^g}{1 - \beta} + \frac{1 - \pi_0}{1 - \beta\lambda} (\bar{v}^b - \bar{v}^g),$$

which is what had to be shown. □

References

- AGHION, P., P. BOLTON, C. HARRIS, and B. JULLIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.
- BANKS, J., and R. SUNDARAM (1998): “Optimal Retention in Principal/Agent Models,” Stern Working Paper FIN-98-006, Stern School of Business, New York University.
- BENSTON, G. (1985): “The Self-Serving Management Hypothesis,” *Journal of Accounting and Economics*, 7, 67–84.
- BROWN, S., and W. GOETZMANN (1995): “Attrition and Mutual Fund Performance,” *Journal of Finance*, 50, 679–698.
- CHEVALIER, J., and G. ELLISON (1998): “Are Some Mutual Fund Managers Better than Others? Cross-Sectional Patterns in Behaviour and Performance,” mimeo.
- FUDENBERG, D., B. HOLMSTROM, and P. MILGROM (1986): “Short-Term Contracts and Long-Term Relationships,” *Journal of Economic Theory*, 51, 1–31.
- GOETZMANN, W., and R. IBBOTSON (1994): “Do Winners Repeat? Patterns in Mutual Fund Performance,” *Journal of Portfolio Management*, 20, 9–17.
- GRUBER, M. (1996): “Another Puzzle: The Growth in Actively Managed Funds,” *Journal of Finance*, 51, 783–810.

- JENSEN, M., and K. MURPHY (1990): “Performance-Pay and Top-Management Incentives,” *Journal of Political Economy*, 98, 225–264.
- KELLER, G., and S. RADY (1999): “Optimal Experimentation in a Changing Environment,” *Review of Economic Studies*, 66(3), 475–507.
- KHORANA, A. (1996): “Top Management Turnover: An Empirical Investigation of Mutual Fund Managers,” *Journal of Financial Economics*, 40, 403–427.
- LAFFONT, J., and J. TIROLE (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 56(5), 1153–1175.
- LAKONISHOK, J., A. SHLEIFER, R. THALER, and R. VISHNY (1991): “Window Dressing by Pension Fund Managers,” *AEA Papers and Proceedings*, 81, 227–231.
- LAMBERT, R. (1986): “Executive Effort and the Selection of Risky Projects,” *Rand Journal of Economics*, 17(1), 77–88.
- OYER, P. (1998): “Fiscal Year Ends and Non-Linear Incentive Contracts: The Effects on Business Seasonality,” *Quarterly Journal of Economics*, 113(1), 149–185.
- RADNER, R. (1985): “Repeated Principal-Agent Games with Discounting,” *Econometrica*, 53(5), 1173–1198.
- RADNER, R. (1986): *Repeated Moral Hazard with Low Discount Rates* Cambridge University Press, Cambridge.
- SCHARFSTEIN, D., and J. STEIN (1990): “Herd Behaviour and Investment,” *American Economic Review*, 80, 465–479.

STOUGHTON, N. (1993): "Moral Hazard and the Portfolio Management Problem," *Journal of Finance*, 48(5), 2009–2028.