

Transparency and Talent Allocation in Money Management

Simon Gervais* Günter Strobl†

Abstract

We construct and analyze a model of delegated portfolio management in which money managers signal their investment skills via their choice of transparency for their fund. We show that a natural equilibrium is one in which high- and low-skill managers pool in opaque funds, while medium-skill managers separate in transparent funds. In this equilibrium, high-skill managers rely on their eventual performance to separate from low-skill managers over time, saving the monitoring costs associated with transparency. In contrast, medium-skill managers rely on transparency to separate from low-skill managers, especially when it is difficult for investors to tell them apart through performance alone. Low-skill managers prefer mimicking high-skill managers in opaque funds in the hope of replicating their performance and compensation. The model yields several novel empirical predictions that contrast transparent funds (e.g., mutual funds) and opaque funds (e.g., hedge funds).

*Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708-0120, sgervais@duke.edu, +1 919 660 7683.

†Frankfurt School of Finance & Management, Sonnemannstrasse 9–11, 60314 Frankfurt, Germany, g.strobl@fs.de, +49 69 154 008 445.

1 Introduction

Starting with Bhattacharya and Pfleiderer (1985), the canonical model of delegated portfolio management has been that of an investor (the principal) hiring a fund manager (the agent) to make investments on his behalf. Several papers within this class of models investigate the possibility that managers signal their skills to investors via the various decisions they make in this context. For example, Huberman and Kandel (1993) and Huddart (1999) show that fund managers increase the risk of their investment strategy in order to signal their ability to gather information that enables them to control and profit from this risk. Similarly, Das and Sundaram (2002) study the joint role of performance-based compensation as an incentive and a signaling device. Finally, Stein (2005) develops a model that explains the prevalence of open-end funds, as such funds do not insure the manager against liquidation risk and so provide a natural signaling avenue for skilled managers.

Our model complements this literature by introducing an important signaling dimension that is available to managers at the inception of their fund: the transparency of the fund. As documented by Almazan et al. (2004), mutual funds differ widely in the constraints that they impose on their manager in an effort to monitor their investment activities and performance. In our model, we assume that transparency allows investors to learn more quickly and more precisely about the skills of the agents they hire to manage their portfolio. That is, investors have access to more than just the historical fund performance to update their beliefs about the manager's ability to invest profitably on their behalf. At the same time, however, we assume that this extra source of information is costly, creating a natural tradeoff: skilled managers want their competence revealed quickly in order to attract more capital to their fund, but take the cost of doing so into account. For example, the decision to operate under the umbrella of a mutual fund family may increase the amount of information that gets communicated to investors about the fund over time (e.g., Gervais, Lynch and Musto, 2005), but producing this information requires time and effort that can erode trading profits.

In this context, we establish that a natural equilibrium for the money management industry is one in which the most and least talented managers make their fund relatively opaque, while managers with moderate abilities make theirs more transparent. The intuition for this result essentially amounts to managers choosing the route that most efficiently facilitates the discovery of their skills by investors. Two instruments are at every manager's disposition: performance and transparency. The fund's performance represents the smallest information set about the manager's skills that emanates from the fund over time. The fund's transparency is chosen at the inception of the fund, and serves to affect the speed at which information about the manager's skill publicly flows to investors.

Because the most skilled managers know that their performance will be difficult to imitate, especially in the long run, they do not see the point of incurring the monitoring costs that come with transparency. That is, they use performance as their main instrument to separate from lesser skilled managers. Managers with more moderate skills face a more subtle tradeoff. They can pool in an opaque fund with the high-skill managers (and the low-skill managers who then have no choice but to also pool) in the hope that they will be able to mimic their performance (and compensation) for a long time. The problem is that the medium-skill managers then also face the possibility that they will be mimicked by low-skill managers, thereby reducing the price that investors are willing to pay for their services. Alternatively, the medium-skill managers can choose to operate a fund that is sufficiently transparent to separate from the low-skill managers who fear the rapid termination of their fund in a transparent environment. This strategy readily conveys their (medium) skill to investors, which is especially beneficial when it is unlikely that they can mimic high-skill managers and likely that they will appear unskilled when pooled with low types. That is, for these medium-skill managers, the immediate convergence of investors' beliefs to their true skill is worth the cost of transparency.

One natural interpretation of this endogenous emergence of opaque and transparent funds in the model is the coexistence of hedge funds and mutual funds. As Brown and Goetzmann (2003) write, "hedge funds are best defined by their freedom from regulatory controls stipulated by the Investment Company Act of 1940" (p. 102). Mutual funds, on the other hand, are heavily regulated and must frequently report detailed information about their positions and strategies to investors. Hedge fund managers are also known to use a wider array of financial instruments than mutual fund managers in order to capitalize on their investment skills (e.g., Fung and Hsieh, 1997; Bookstaber, 2003). In contrast, Koski and Pontiff (1999) document that mutual fund managers make little use of complex securities in their trading.

In this light, the aforementioned partial-pooling equilibrium suggests that hedge funds will tend to be run by managers with extreme skills, high and low, whereas mutual funds will tend to attract managers with more average talent. The model's endogenous fund flows, excess returns, and manager compensation then lead to several empirical implications contrasting mutual funds and hedge funds. For example, we expect the cross-sectional dispersion of fund alphas to be greater for hedge funds than for mutual funds. We also expect hedge fund flows to be more sensitive to past performance than mutual fund flows.

Brown et al. (2008) empirically investigate the relationship between the disclosures of hedge funds and their operational risk. One of their findings is that hedge funds are more likely to disclose information about their fund following good performance, leading the authors to conclude that filing may be a signal of quality. Viewed through the lens of our model, their results may

indicate that the early lack of filing is used by managers to signal their skills which subsequently get revealed through the good performance of their fund.

Our model builds on Berk and Green's (2004) insight that capital should rapidly flow to funds that experience high returns, as investors update about the fund manager's skill for investing profitably. In equilibrium, the marginal benefit of investing in the fund is equal to the marginal cost, funds do not generate any excess (net) returns, and the entire economic surplus accrues to the manager in the form of compensation for his services. Like Dangl, Wu and Zechner (2008), we add the possibility that the fund controls the speed at which investors learn the manager's skills. In their model, this is done privately by the fund's management company which strategically times the replacement of the manager in order to positively influence the updating process of investors about the quality of the fund. In our model, this is done publicly by the manager himself who self-selects into a transparency level that subsequently impacts the updating process of investors. As such, the choice of transparency acts as a signal in our model, whereas it creates a moral hazard problem in theirs.

Our results relate to a recent paper by Daley and Green (2011). They show that the addition of independent signals about skills (which they call "grades") in Spence's (1973) seminal model of job-market signaling makes high-skill and low-skill workers more likely to pool in equilibrium. The analogy to grades in our model is the fund manager's performance, which cannot be controlled. When the performance of high types cannot be easily imitated, they are more than happy to rely exclusively on their performance in order to separate from other fund managers in the long run. That is, they do not mind pooling in the short run, as their performance will eventually speak for itself. A similar result is obtained by Alós-Ferrer and Prat (2011) in a job-market signaling model with dynamic learning. As in our model, the receiver (employer) learns about the sender's (worker's) type over time and, when learning is fast and accurate, a pooling equilibrium emerges. Our paper differs from these two papers by considering a third (medium) type who is not as skilled as the high type but, unlike the low type, can create some value.

This third type makes our results more similar to those obtained by Feltovich, Harbaugh and To (2002) in the context of job-market signaling with three worker types. As we show, the partial-pooling equilibrium that we derive is often the Pareto-dominant equilibrium, in that it makes all money managers better off. This is due to the fact that, in a money-management context, the cost of signaling through higher levels of transparency is (endogenously) higher for high types as transparency leads to valuable information leakages and thus erodes more value for those managers who have the potential to generate consistently good performance. That is, for highly skilled managers, separation is cheaper to achieve via long-run performance than via the extra information about their type that is facilitated by making their fund transparent.

The rest of the paper proceeds as follows. In section 2, we introduce the model. Section 3 contains our equilibrium analysis. The empirical predictions of the model are derived and discussed in section 4. Finally, section 5 summarizes and concludes. All proofs are contained in the Appendix.

2 The Model

We consider an N -period economy populated with risk-neutral money managers and investors who discount future cash flows at a rate normalized to zero. Money managers (or just *managers* for short) have no wealth, but potentially have some investment skills. Investors are wealthy, but have no investment skills. Each money manager can open an investment fund (or just *fund* for short), and investors must decide on how much capital A_n to allocate to each fund at the beginning of every period $n \in \{1, \dots, N\}$. Money managers come in three different types, high, medium or low, that correspond to their ability to generate investment profits. Specifically, the manager's type is drawn from the following distribution:

$$\tilde{\tau} = \begin{cases} h, & \text{prob. } \lambda_h \\ m, & \text{prob. } \lambda_m \\ \ell, & \text{prob. } \lambda_\ell, \end{cases} \quad (1)$$

where $\lambda_h + \lambda_m + \lambda_\ell = 1$. The manager privately observes his own type at the outset. This type determines the distribution of the n^{th} -period excess return $\tilde{r}_n(\tilde{\tau})$ that his fund will generate. We assume that these excess returns (or just *returns* for short) are independently drawn across periods. The performance of high types first-order stochastically dominates that of medium types, which in turn first-order stochastically dominates that of low types. Although any first-order stochastic ordering is sufficient to generate our results, the analysis is greatly simplified by assuming three-point distributions (good, average, and bad) for the various types' returns. Specifically, for the low types, we assume that

$$\tilde{r}_n(\ell) = \begin{cases} r_G, & \text{prob. } p_G \\ r_A, & \text{prob. } p_A \\ r_B, & \text{prob. } p_B, \end{cases} \quad (2)$$

where $p_G + p_A + p_B = 1$, $r_G > r_A > r_B$, and $\mu_\ell \equiv p_G r_G + p_A r_A + p_B r_B = 0$. That is, on average, the low types do not generate any excess returns, and so investors do not benefit from their investment services.¹ The medium types' returns in period n are given by

$$\tilde{r}_n(m) = \begin{cases} r_G, & \text{prob. } \frac{p_G}{p_G + p_A} \\ r_A, & \text{prob. } \frac{p_A}{p_G + p_A}, \end{cases} \quad (3)$$

and their expected returns are $\mu_m \equiv \frac{p_G r_G + p_A r_A}{p_G + p_A} > 0$. Finally, the high types' return in period n are assumed to be equal to $r_G \equiv \mu_h > \mu_m$ with probability one. These distributions, although

¹Our results are unaffected if we assume more generally that $\mu_\ell \leq 0$.

somewhat specific, capture the notion that the one-period performance of low and medium types will match that of high types with some probability (p_G and $\frac{p_G}{p_G+p_A}$, respectively), and the low types' performance will match that of medium types with probability $p_G + p_A$. The advantage, which will become more apparent later, comes from the fact that skilled managers can never be mistaken for less skilled managers, which greatly simplifies the analysis.

At the outset, managers choose the *transparency* $t \in [0, 1]$ of their fund. This choice, which is publicly observable by investors, can be used by managers to signal their type to investors in equilibrium. A fund's transparency is meant to capture the possibility that money management vehicles can take different forms, and our assumption is that these various forms are indexed by their transparency. For example, hedge funds are known to be more opaque about their operations than mutual funds, and some hedge funds are less open to divulging their strategies to potential investors than other hedge funds. This, we assume, translates into various degrees of screening by investors: more transparent funds allow investors to better sort managers with respect to their skills. That is, the possibility of looking at more than just the return history of a fund allows investors to determine with more accuracy whether a manager's performance is due to his skill or to (bad) luck.

In terms of the model, we assume that a more transparent fund allows investors to identify early the money managers who do not have the ability to generate any excess returns (i.e., the low types). Specifically, we assume that, before the fund is open for potential investment, investors observe a signal $\tilde{u}_t \in \{0, 1\}$, which has the following distribution:

$$\Pr\{\tilde{u}_t = 0 \mid \tilde{\tau} = \ell\} = t = 1 - \Pr\{\tilde{u}_t = 1 \mid \tilde{\tau} = \ell\}, \quad (4)$$

$$\Pr\{\tilde{u}_t = 1 \mid \tilde{\tau} = m\} = \Pr\{\tilde{u}_t = 1 \mid \tilde{\tau} = h\} = 1. \quad (5)$$

In effect, therefore, a signal $\tilde{u}_t = 0$ makes it clear to investors that the manager's type is low and that they should not hire him, as such a signal can never be observed when the manager's type is medium or high. Conversely, a signal $\tilde{u}_t = 1$ reduces the possibility that the manager they are about to hire is a low type, as

$$\Pr\{\tilde{\tau} = \ell \mid \tilde{u}_t = 1\} = \frac{\Pr\{\tilde{u}_t = 1 \mid \tilde{\tau} = \ell\} \Pr\{\tilde{\tau} = \ell\}}{\sum_{\tau \in \{h, m, \ell\}} \Pr\{\tilde{u}_t = 1 \mid \tilde{\tau} = \tau\} \Pr\{\tilde{\tau} = \tau\}} = \frac{(1-t)\lambda_\ell}{\lambda_h + \lambda_m + (1-t)\lambda_\ell}$$

is decreasing in t . Implicitly, but somewhat deliberately, the assumed distribution for \tilde{u}_t implies that investors cannot distinguish between high types and medium types. That is, high types do not have an obvious advantage over medium types about signaling their skill; investors can only tell whether a manager's strategy has potential but they cannot tell how profitable it will be. The model's results hold as long as transparency allows investors to disentangle low types from other

types more accurately than medium types from high types. Less critical is the fact that investors do not get to observe an independent draw from \tilde{v}_t 's distribution in every period. This possibility would certainly allow investors to converge more rapidly on the true type of the manager, but would not affect the main economic forces of the model. The gain in tractability from our one-shot signal assumption is thus warranted.

Since more transparency can only eliminate low types from the pool of money managers, medium and high types would automatically choose $t = 1$ if this choice were not costly. To add some tension to this choice, we assume that more transparency comes with a higher per-dollar cost of managing a fund's assets. Following Berk and Green (2004) and in light of Chen et al.'s (2004) empirical findings, we assume that the per-dollar cost of managing a fund is increasing in the size of the fund. In fact, to keep things simple without any loss of generality, we assume that the per-dollar cost of managing a fund of size A_n in period n is $k_t A_n$, where $k_t > 0$. Our innovation over Berk and Green (2004) is that we assume that k_t is strictly increasing in t , capturing the idea that a more transparent fund requires more costly information to be produced, more constraints on potential investments to be made (e.g., no short-selling or no investment in privately held firms), and potentially leads to leakages of strategies that reduce their effective profitability.

The manager announces his compensation contract at the beginning of each period n . This contract is represented by the amount w_n that the manager charges to investors per dollar they invest in the fund at the beginning of the period.² Based on this contract, on the history of the fund's returns, on \tilde{v}_t , and on rational equilibrium updates about the manager's type, investors choose the amount A_n of money that they allocate to the fund for this one period. Let us denote the period- n profits of investors by $\tilde{\pi}_n \equiv A_n [\tilde{r}_n(\tilde{\tau}) - w_n - k_t A_n]$. As in Berk and Green's (2004) model, we assume that investors compete for the manager's scarce talent for investing, and so the equilibrium A_n always makes them indifferent between investing in the fund and not doing so; that is, in equilibrium, investors make zero expected profits and managers capture the entire economic surplus that they create.

Lemma 1. *Suppose that the investors' information set is \mathcal{I}_n at the beginning of period n . The amount they invest in the fund in that period is then*

$$A_n = \max \left\{ 0, \frac{\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n] - w_n}{k_t} \right\}. \quad (6)$$

²Note that the period- n contract could be made contingent on the return $\tilde{r}_n(\tilde{\tau})$ of the fund in that period. Given that our model abstracts from moral hazard problems, this would have no effect on fund performance in that period. It would however allow a skilled manager to capture more of the available surplus than a less skilled manager with the same history, as in Gervais, Lynch and Musto (2005). Since this consideration does not affect our results, we stick with the simpler flat per-period wage.

The max operator in (6) simply says that investors are not willing to invest any money when the manager's compensation exceeds the excess returns that his fund is expected to generate. In fact, we assume that the fund closes at that point; that is, the sequence of observable returns $\tilde{r}_n(\tilde{\tau})$ ceases as soon as the manager cannot attract any money into his fund. In this sense, (6) also highlights the fact that the manager cannot profitably manage a fund once investors reach an information set \mathcal{I}_n such that $\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n] \leq 0$.

At the beginning of every period n , the manager chooses his contract in order to maximize his expected compensation, $U_n \equiv w_n A_n$, for that period. In doing so, he takes the investors' reaction to his choice, (6), into account. The following lemma characterizes his decision.

Lemma 2. *Suppose that the investors' information set is \mathcal{I}_n at the beginning of period n , and that*

$$\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n] > 0. \quad (7)$$

Then the manager's choice of per-dollar compensation for that period is

$$w_n = \frac{1}{2} \mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n], \quad (8)$$

the size of his fund is

$$A_n = \frac{1}{2k_t} \mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n], \quad (9)$$

and his total compensation is

$$U_n = w_n A_n = \frac{1}{4k_t} (\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n])^2. \quad (10)$$

Note that the manager can anticipate the investors' beliefs and the corresponding amount that they will be willing to invest for every return sequence that his fund generates. Thus, the assumption that he announces his contract only at the beginning of each period is equivalent to him announcing a set of path-dependent contracts for every period $n \in \{1, \dots, N\}$ at the inception of the fund. From (9) and (10), we can also see that our assumption that transparency affects the per-dollar cost of the fund is equivalent to the alternative assumption that the returns of a fund with transparency t are scaled by a constant of $\frac{1}{k_t}$. That is, because investors break even in every period, it does not matter if the costs of transparency are felt through lower gross fund returns (through strategy leaks, for example) or through more expensive administrative costs.³

Lemma 2 shows that the per-dollar compensation of a manager is proportional to his fund's expected return, conditional on the investors' information set. Since a transparency level of t

³The empirical evidence about the effect of transparency on fund performance appears to be mixed. For example, while Frank et al. (2004) find that mutual fund disclosures tend to attract investment strategy imitators and reduce gross performance, Aggarwal and Jorion (2011) find that more transparent hedge funds do not seem to under-perform (and perhaps outperform) their more opaque counterparts.

eliminates a fraction t of low-type managers, this implies that the ex ante expected compensation of a manager of type $\tilde{\tau} = \ell$ is proportional to $(1-t)\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n]$. As we show later, this quantity is always decreasing in t , despite the fact that $\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n]$ is increasing in t . However, since the fund size A_n is also proportional to $\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n]$, the total compensation that a low type can expect ex ante is actually a quadratic function of his fund's expected return and, hence, can be increasing in t . To rule out the economically less sensible case in which all manager types benefit from more transparency, we assume that the cost function k_t increases sufficiently fast in t .

Assumption 1. $\frac{dk_t}{dt} \geq \frac{\lambda_\ell}{\lambda_h} k_t$, for all $t \in [0, 1]$.

As will become clear in section 3, this assumption ensures that the fund size A_n , given by (9), does not increase in t , which further implies that the expected total compensation of a low-type manager decreases in t .

To summarize, the timeline for the game is as follows. At the outset, the manager privately observes his type $\tilde{\tau}$. Upon learning this type, he opens a fund by publicly announcing a transparency level t for it. Before any investment is made into the fund, the investors observe \tilde{i}_t . In the event that $\tilde{i}_t = 0$, the game ends as the manager is readily identified as a low type. If instead $\tilde{i}_t = 1$, then N investment periods follow. Each of these periods involves the same subsequence of events. First, the manager announces his per-dollar compensation w_n . Second, the investors choose the amount A_n of capital to invest in the fund. Third, the manager receives his compensation, the fund return for the period is realized, and investors receive their net profits.⁴ The game ends (i.e., the fund closes) when the investors' information set \mathcal{I}_n at the beginning of any period n indicates that $\mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n] \leq 0$, or at the end of N periods, whichever comes first.

3 Equilibrium Analysis

Because the manager announces his fund's transparency after he learns his type, his publicly observable choice of t can therefore be used to signal his private type $\tilde{\tau}$. To keep the analysis intuitive, we only consider pure-strategy equilibria of this signaling game. In any such equilibrium, the investors update their beliefs about the manager's type based on the equilibrium strategy of money managers according to Bayes' rule, and money managers cannot profitably deviate from the equilibrium by choosing a different transparency level.

As is typical in signaling games, there are multiple equilibria to this game, and so some subjectivity remains in selecting the one(s) that should prevail. Our equilibrium analysis proceeds as follows. First, we eliminate some equilibria that can never exist. Then we characterize the only

⁴The model is unaffected if we assume that the capital invested at the beginning of the period is returned to these investors at the end of the period, before a new round of capital is raised in the subsequent period.

two types of equilibria that potentially exist. Finally, we identify intuitive conditions that make one such equilibrium Pareto-dominant. Before we embark on this analysis, it is worth noting that few results can be borrowed from standard game theory as the costs of signaling in our model are non-monotonic in types. This happens for two reasons. First, an increase in transparency not only increases the per-dollar cost $k_t A_n$ of a fund, but it also creates the possibility (for low types) of losing one's job (at the outset if $\tilde{\tau}_t = 0$). Second, although k_t does not depend on the manager's type, the cost of making a fund more transparent is effectively larger for high types. Indeed, over time, the reputation of high types tends to be higher than that of other types and as a result their funds tend to be better capitalized. The per-period cost of transparency, $k_t A_n^2$, therefore tends to be larger for high types.

3.1 Preliminaries

Let us denote the strategy adopted by each of the three types in a given game by the transparency they choose for their fund: $\mathcal{S} = \{t_h, t_m, t_\ell\} \in [0, 1]^3$. The following proposition restricts the set of potential equilibria for the game.

Proposition 1. *The equilibrium $\{t_h, t_m, t_\ell\}$ to the signaling game must have $t_h = t_\ell$.*

This result derives from the fact that, in equilibrium, it must always be the case that high types pool with some other type, and that low types pool with some type as well. It is straightforward to verify that at least one of these two conditions fails to hold when $t_h \neq t_\ell$.

To see the intuition behind each of these conditions, suppose first that the high types separate using a transparency level of t_h that is different from both t_m and t_ℓ . Let us denote their expected N -period compensation from employing this strategy by $u_h(t_h)$, where $u_\tau(t) = \mathbb{E}[U(t) \mid \tilde{\tau} = \tau]$ and $U(t) = \sum_{n=1}^N U_n(t)$. Since the high types can always imitate the medium types without ever risking being misperceived as a low type (i.e., $\Pr\{\tilde{\tau}_n(h) = r_B\} = 0$), we must have $u_h(t_h) > u_h(t_m) \geq u_m(t_m)$, where the last inequality derives from the fact that high types can never do worse than the medium types if they use the same strategy (since their returns first-order stochastically dominate those of medium types). But since investors readily conclude that transparency t_h identifies a high type, they do not update their beliefs about the manager whose fund transparency is t_h .⁵ Thus the medium types are immediately tempted to imitate them as $u_m(t_h) = u_h(t_h) > u_m(t_m)$, breaking the equilibrium in the process. The second condition, that low types always pool, is even more intuitive. By separating, low types identify themselves as managers who cannot generate excess

⁵The critical reader might argue that this result is inconsistent with the fact that the support of the return distribution of high types differs from that of medium types. We want to point out, however, that introducing an arbitrarily small probability that high types generate a return of r_A would completely eliminate this apparent inconsistency.

returns, and so they can never attract any capital from investors. They are therefore always strictly better off when they pool with one of the other types, as this allows them to at least open a fund.

3.2 Partial-Pooling Equilibrium

By Proposition 1, any equilibrium must have high types and low types pooling, but it can be the case that medium types separate. In this and the next subsection, we explore the two possible equilibria, starting with the partial-pooling equilibrium $\{t, t', t\}$ where $t' \neq t$.

Let $\phi_1(t)$ denote the initial fraction of high types amongst the managers who choose transparency t . Since the signal \tilde{i}_t eliminates a fraction t of the low-type managers, we have

$$\phi_1(t) = \frac{\lambda_h}{\lambda_h + \lambda_\ell(1-t)}. \quad (11)$$

As long as such a fund generates a return of r_G , investors will keep allocating capital to the manager. In fact, with every return of r_G , investors increase their beliefs that the manager is a high type. Conversely, they (correctly) conclude that he is a low type as soon as the return in any given period is r_A or r_B . In that event, the fund closes as investors never invest with a manager who cannot generate excess returns. Our next result characterizes the updating process of investors for fund t . To derive it, it is useful to introduce some notation. Specifically, let $\tilde{r}^n \in \{r_G, r_A, r_B\}^n$ denote a return sequence over the fund's first n periods, and let r_G^n denote the return sequence with returns of r_G in each of the first n periods.

Lemma 3. *In an equilibrium $\{t, t', t\}$, the investors' beliefs about the type of the manager of fund t at the beginning of period n are given by*

$$\Pr\{\tilde{\tau} = h \mid \tilde{i}_t = 1, \tilde{r}^{n-1} = r_G^{n-1}\} = \frac{\lambda_h}{\lambda_h + \lambda_\ell(1-t)p_G^{n-1}} \equiv \phi_n(t), \quad (12)$$

$$\text{and } \Pr\{\tilde{\tau} = h \mid \tilde{i}_t = 1, \tilde{r}^{n-1} \neq r_G^{n-1}\} = 0. \quad (13)$$

Notice that $\phi_n(t)$ converges to one as n increases to infinity: investors become increasingly confident that the manager is a high type as they keep experiencing high returns in his fund. In fact, the results of Lemma 3 can be combined with those of Lemma 2 to calculate the returns that investors can expect from fund t over time, as well as the ex ante expected compensation of the (high type and low type) managers of such funds.

Lemma 4. *In an equilibrium $\{t, t', t\}$, the returns that investors expect from fund t at the beginning of period n are*

$$\mathbb{E}[\tilde{r}_n(\tilde{\tau}) \mid \tilde{i}_t = 1, \tilde{r}^{n-1} = r_G^{n-1}] = \phi_n(t)\mu_h = \frac{\lambda_h\mu_h}{\lambda_h + \lambda_\ell(1-t)p_G^{n-1}} \equiv \bar{r}_n(t), \quad (14)$$

$$\text{and } \mathbb{E}[\tilde{r}_n(\tilde{\tau}) \mid \tilde{i}_t = 1, \tilde{r}^{n-1} \neq r_G^{n-1}] = 0. \quad (15)$$

The fund closes at the end of the period n for which a return other than r_G is observed for the first time. Upon the manager learning his type, the expected N -period compensation of a high-type manager is

$$u_h(t) = \frac{1}{4k_t} \sum_{n=1}^N [\bar{r}_n(t)]^2, \quad (16)$$

and that of a low-type manager is

$$u_\ell(t) = \frac{1-t}{4k_t} \sum_{n=1}^N p_G^{n-1} [\bar{r}_n(t)]^2. \quad (17)$$

Naturally, the returns that investors expect from the fund keep rising as they put more and more weight on the possibility that the manager is a high type. This can be seen from (14), which converges to μ_h as $n \rightarrow \infty$. Of course, this happens only when the sequence of returns generated by the fund is consistently high (i.e., every period's return is equal to r_G). Since this happens with probability one for the high type, he can expect to receive $\frac{1}{4k_t} [\bar{r}_n(t)]^2$ with certainty in each period n , as seen in (16).

On the other hand, the low-type manager collects this compensation only in the event that he mimics the performance of a high-type manager. As reflected in (17), this means two things. First, the low type has to survive the initial round of monitoring with a signal $\tilde{v}_t = 1$; this happens with probability $1 - t$. Second, his fund has to generate a return of r_G in every period and, for him, this happens with probability p_G in every period. As a result, the compensation that he can expect to collect upon learning his type comes predominantly from the early periods, as it is unlikely that he will be able to replicate a high type's performance in the long run.

Let us now turn our attention to the medium types whose funds have a transparency level $t' \neq t$. Since these managers separate, investors readily conclude that a fund with a transparency level t' is managed by a medium type. That is, they do not update: they expect returns of μ_m in every period, and they rationally ignore the fund's return sequence.

Lemma 5. *In an equilibrium $\{t, t', t\}$, the returns that investors expect from fund t' at the beginning of period n is $\mu_m > 0$, and the fund never closes. Upon the manager learning his type, the expected N -period compensation of a medium-type manager is*

$$u_m(t') = \frac{1}{4k_{t'}} \sum_{n=1}^N \mu_m^2 = \frac{N\mu_m^2}{4k_{t'}}. \quad (18)$$

Since $\mu_m < \mu_h$, medium types do not generate expected returns that are as high as the ones that high types generate. Still, because $\mu_m > 0$, medium types create some value for investors. So, when these managers are identifiable without any uncertainty (through their choice of transparency

$t' \neq t$), investors are willing to capitalize their fund in every period, regardless of the fund's performance over time. That is, in equilibrium, investors assign a 100% probability to the event that the manager is a medium type and so need not look at his performance for more information.

We now proceed to derive the conditions under which the conjectured partial-pooling equilibrium exists. For $\{t, t', t\}$ to be an equilibrium, it must be the case that the high and low types do not benefit from choosing t' instead of t for their fund, and vice versa for the medium types. Further, no type can be better off choosing an out-of-equilibrium transparency level $\hat{t} \notin \{t, t'\}$. This latter condition will be satisfied whenever the investors' out-of-equilibrium beliefs assign a sufficiently high probability to the event that the deviating manager is of type $\tilde{\tau} = \ell$.

Let us start with the low types. For these managers, choosing t' instead of t means two things. First, it means that they subject themselves to a different level of scrutiny before they can even open their fund. In particular, if $t' > t$, then choosing t' makes it less likely that they will even be able to get their fund off the ground as they run a higher risk of getting discovered as low types. Second, if \tilde{t}' turns out to be one, investors assign a 100% probability to the event that they are a medium type and so the managers' expected N -period compensation at that point is equal to that of medium types, as calculated in (18).⁶ Thus, at the time of choosing his fund's transparency, the deviating low type's ex ante expected N -period compensation is equal to

$$\hat{u}_\ell(t') = (1 - t') \frac{N\mu_m^2}{4k_{t'}}. \quad (19)$$

For $\{t, t', t\}$ to be an equilibrium therefore, it must be the case that t' is sufficiently large to make (19) smaller or equal to (17). Since $\hat{u}_\ell(t')$ is strictly decreasing in t' and is equal to zero at $t' = 1$, there is always a range $[\bar{t}(t), 1]$ for t' that makes this possible, where $\bar{t}(t)$ is the transparency level t' that makes (19) equal to (17). Intuitively, no matter how attractive the prospect of managing under the guise of a medium type is to a low type, there is always a level of transparency that can be adopted by the medium types to separate from the low types.

By deviating to t , the medium types can expect a total compensation of $\frac{1}{4k_t} [\bar{r}_n(t)]^2$ in period n , as derived in Lemma 4, but only as long as they generate returns of r_G . Since, for them, this happens with probability $\frac{p_G}{p_G + p_A}$ in each period, their expected N -period compensation is then equal to

$$\hat{u}_m(t) = \frac{1}{4k_t} \sum_{n=1}^N \left(\frac{p_G}{p_G + p_A} \right)^{n-1} [\bar{r}_n(t)]^2. \quad (20)$$

⁶Note that investors do not update their beliefs if a low type chooses an out-of-equilibrium transparency level t' even though this manager will sometimes generate a return of r_B . As argued above, this apparent inconsistency is due to our assumption that the return distributions for the three manager types do not share the same support. Allowing for an arbitrarily small probability that the medium type generates a return of r_B would completely resolve this issue.

A comparison of (18) and (20) reveals that the medium types do not benefit from deviating for N sufficiently large, as $\left(\frac{p_G}{p_G+p_A}\right)^{n-1} [\bar{r}_n(t)]^2$ decreases to zero as $n \rightarrow \infty$ and so is eventually lower than $\frac{1}{4k_{t'}}\mu_m^2$. Intuitively, the medium types do not benefit from imitating high types and putting their long-term compensation at risk when they can rely on a long stream of sure compensation by separating from them.

Finally, consider a high type who deviates to t' . Since, in equilibrium, he is immediately identified as a medium type, his N -period expected compensation from doing so is equal to that of medium types, as derived in Lemma 5:

$$\hat{u}_h(t') = \frac{1}{4k_{t'}} \sum_{n=1}^N \mu_m^2 = \frac{N\mu_m^2}{4k_{t'}}. \quad (21)$$

Comparing this quantity to (16), we see that the deviation will not be beneficial for the high type if N is sufficiently large: since $\bar{r}_n(t)$ converges to $\mu_h > \mu_m$ as $n \rightarrow \infty$, the compensation that he can expect in the long run is larger with a transparency level t rather than t' as long as $t < t'$.⁷ The following proposition summarizes this discussion.

Proposition 2. *A partial-pooling equilibrium $\{t, t', t\}$ with $t < t' \in [\bar{t}(t), 1]$ exists provided that N is sufficiently large.*

The partial-pooling equilibrium described in Proposition 2 has the medium types choose a higher degree of transparency for their fund than the high and low types. Intuitively, in this equilibrium, the medium types use transparency (i.e., a large t') to separate from the low types, whereas the high types use a combination of transparency (t) and long-run performance to eventually achieve separation from low types.

As is common in signaling models, our model admits multiple perfect Bayesian equilibria. In fact, the indeterminacy of investor beliefs off the equilibrium path gives rise to a continuum of partial-pooling equilibria with different transparency levels. However, in the following we argue—and formally establish by applying the refinement concept proposed by Mailath, Okuno-Fujiwara, and Postlewaite (1993)—that, for N large enough, the only plausible partial-pooling equilibrium is the Pareto-dominant one with the lowest level of transparency.

Since higher levels of transparency come with a higher per-period cost of managing a fund, it is worth asking why the high types should pick a $t > 0$ in the first place. The answer is that, in equilibrium, this eliminates a fraction $1 - t$ of the low types who seek to imitate them with the same transparency level. The benefit is short-lived, however, as the low types get eliminated through time when their performance is not on par with that expected from high types (i.e., when their

⁷For a more formal argument, see the proof of Proposition 2.

fund experiences a return of r_A or r_B). More technically, while $\bar{r}_1(t)$ can be much greater than $\bar{r}_1(0)$, the difference between $\bar{r}_n(t)$ and $\bar{r}_n(0)$ is negligible for large n , as both quantities converge to μ_h as $n \rightarrow \infty$. For the high types therefore, any transparency level $t > 0$ eventually leads to a waste of surplus in the form of extra per-period management costs. That is, as $n \rightarrow \infty$, their expected total compensation in period n converges to $\frac{\mu_h^2}{4k_t}$ when it could converge to a higher value of $\frac{\mu_h^2}{4k_0}$ with $t = 0$. Thus, amongst the various partial-pooling equilibria of Proposition 2, high types strictly prefer those with $t = 0$.

Similarly, for any given t chosen by the high types, the medium types prefer the lowest possible level of transparency that allows them to separate from the low types. That is, they prefer $t' = \bar{t}(t)$ to any higher level of transparency, as they prefer a per-period compensation of $\frac{\mu_m^2}{4k_{\bar{t}(t)}}$ to $\frac{\mu_m^2}{4k_{t'}}$ for all $t' \in (\bar{t}(t), 1]$. Interestingly, the high types' preference for $t = 0$ can benefit the low types as well. Indeed, for the low types, the option to manage a fund without any monitoring allows them to capture at least one period of compensation with certainty by pooling with the high types.

This intuition can be formalized with the refinement of *undefeated equilibria* introduced by Mailath, Okuno-Fujiwara, and Postlewaite (1993).⁸ This refinement is appealing because it requires out-of-equilibrium beliefs to be “globally” consistent, in the sense that any adjustment of out-of-equilibrium beliefs has to be consistent with beliefs at other information sets, including those along the equilibrium path.⁹ We present a formal definition of the refinement in the Appendix. Intuitively, it works as follows. Consider an equilibrium \mathcal{S} and a transparency level t' that is not chosen in the equilibrium. Suppose there is an alternative equilibrium \mathcal{S}' in which some non-empty set of types of fund manager choose transparency t' and that this set is precisely the set of types who prefer the alternative equilibrium \mathcal{S}' to the original equilibrium \mathcal{S} . Then the refinement requires the investors' beliefs following t' in equilibrium \mathcal{S} to be consistent with this set. If they are not, the alternative equilibrium \mathcal{S}' *defeats* equilibrium \mathcal{S} . The following proposition shows that the refinement eliminates all partial-pooling equilibria except for the Pareto-dominant one.

Proposition 3. *For sufficiently large N , the only undefeated partial-pooling equilibrium is the equilibrium $\{0, \bar{t}(0), 0\}$.*

In the equilibrium of Proposition 3, high types and low types manage fully opaque funds, while medium types manage funds that are transparent enough to keep low types away. This equilibrium is reminiscent of Feltovich, Harbaugh and To's (2002) partial-pooling equilibrium, in which high-type workers refrain from using schooling as a costly signaling device to separate from low-type

⁸For other applications of this refinement see, for example, Spiegel and Spulber (1997), Taylor (1999), Gomes (2000), and Fishman and Hagerty (2003).

⁹This property distinguishes the refinement of undefeated equilibria from other belief-based refinements such as the intuitive criterion (Cho and Kreps, 1987) and perfect sequential equilibrium (Grossman and Perry, 1986).

workers (as they do in Spence (1973)) when an independent source of information that they do not control is likely to reveal their high skill to employers.¹⁰ In their equilibrium, the medium-type workers prefer separating from low-type workers, as they fear the likely possibility that the independent source of information will not identify them away from low-type workers.

In our model, transparency is the costly signal that medium types use to separate from the low types, as they fear the likely possibility that their performance, the independent source of information that is imperfectly correlated with their type and that they do not control, will not allow investors to quickly identify them away from low types. High types, however, prefer saving on signaling costs with a low transparency level that minimizes monitoring costs, as they can be confident that their performance will be drastically different from that of low types, allowing investors to quickly compensate them as high types.

3.3 Pooling Equilibrium

As shown in Proposition 1, partial-pooling equilibria are not the only equilibria that can prevail. In particular, although this proposition shows that high types and low types always choose the same transparency for their fund, it is possible that medium types also refrain from separating. In this section, we characterize such pooling equilibria.

When all types pool with a transparency of t , investors can only use the signal \tilde{v}_t and the return sequence of the fund to update their beliefs about the manager's type. Given the return process $\tilde{r}_n(\tilde{\tau})$ specified in section 2, investors can eliminate the possibility that the manager is a high type as soon as they observe a return of r_A or r_B , and readily conclude that the manager is a low type as soon as they observe a return of r_B . The following lemma characterizes the expected returns that result from this belief-updating process by investors. In addition to r_G^n introduced earlier, it uses \mathcal{R}_A^n to denote the subset of all n -period return sequences that do not include a return of r_B .

Lemma 6. *In a pooling equilibrium $\{t, t, t\}$, the returns that investors expect at the beginning of period n are given by*

$$\mathbb{E}[\tilde{r}_n(\tilde{\tau}) \mid \tilde{v}_t = 1, \tilde{r}^{n-1} = r_G^{n-1}] = \frac{\lambda_h \mu_h + \lambda_m \left(\frac{p_G}{p_G + p_A} \right)^{n-1} \mu_m}{\lambda_h + \lambda_m \left(\frac{p_G}{p_G + p_A} \right)^{n-1} + \lambda_\ell (1-t) p_G^{n-1}} \equiv \bar{\rho}_n(t), \quad (22)$$

$$\mathbb{E}[\tilde{r}_n(\tilde{\tau}) \mid \tilde{v}_t = 1, \tilde{r}^{n-1} \in \mathcal{R}_A^{n-1}, \tilde{r}^{n-1} \neq r_G^{n-1}] = \frac{\lambda_m \mu_m}{\lambda_m + \lambda_\ell (1-t) (p_G + p_A)^{n-1}} \equiv \bar{\bar{\rho}}_n(t), \quad (23)$$

$$\text{and } \mathbb{E}[\tilde{r}_n(\tilde{\tau}) \mid \tilde{v}_t = 1, \tilde{r}^{n-1} \notin \mathcal{R}_A^{n-1}] = 0. \quad (24)$$

¹⁰Daley and Green (2011) refer to this independent source of information as a “grade,” and show that more pooling occurs even when there are only two types of workers.

Since $\bar{\rho}_n(t)$ converges to μ_h as $n \rightarrow \infty$, the high types will eventually get credit for the returns that they generate, as their funds always return r_G . Similarly, the funds run by medium types will eventually be expected to generate μ_m , as $\bar{\rho}_n(t)$ converges to μ_m and medium types never generate returns below r_A . Finally, although the low types can imitate high and low types for a while, they are eventually found to be low types with their first return of r_B and must close their fund at that point.

The period- n compensation that each type can expect at the outset therefore depends on the probability that their fund looks like that of a high type or that of a medium type at the beginning of period n . The following lemma derives the expected N -period compensation that a manager can expect upon learning his type at the outset.

Lemma 7. *In a pooling equilibrium $\{t, t, t\}$, the expected N -period compensation of a high-type manager is*

$$u_h(t) = \frac{1}{4k_t} \sum_{n=1}^N [\bar{\rho}_n(t)]^2, \quad (25)$$

that of a medium-type manager is

$$u_m(t) = \frac{1}{4k_t} \sum_{n=1}^N \left\{ [\bar{\rho}_n(t)]^2 + \left(\frac{p_G}{p_G + p_A} \right)^{n-1} \left([\bar{\rho}_n(t)]^2 - [\bar{\rho}_n(t)]^2 \right) \right\}, \quad (26)$$

and that of a low-type manager is

$$u_\ell(t) = \frac{1-t}{4k_t} \sum_{n=1}^N \left\{ (p_G + p_A)^{n-1} [\bar{\rho}_n(t)]^2 + p_G^{n-1} \left([\bar{\rho}_n(t)]^2 - [\bar{\rho}_n(t)]^2 \right) \right\}. \quad (27)$$

Since $\lim_{n \rightarrow \infty} \bar{\rho}_n(t) = \mu_m$ and $\lim_{n \rightarrow \infty} \left(\frac{p_G}{p_G + p_A} \right)^{n-1} = 0$, we can see from (26) that the medium types expect to receive $\frac{\mu_m^2}{4k_t}$ per period in the long run. Before that, since $\bar{\rho}_n(t) > \bar{\rho}_n(t)$, they extract some rent from pooling with high types as long as they can imitate their performance. Similarly, the low types who survive the initial monitoring round (i.e., with $\tilde{u}_t = 1$) can expect to receive zero per period in the long run and their fund to eventually close. Before that, a low type extracts some rents from pooling with higher types as long as his performance mimics theirs.

For the same reason that high types prefer a fully opaque fund (i.e., a transparency level of zero) amongst all of the partial-pooling equilibria, they and medium types also prefer the pooling equilibrium $\{t, t, t\}$ that has $t = 0$. Indeed, in the long run, they both prefer the lower monitoring costs, larger funds, and larger compensation that come with $t = 0$. It is therefore not surprising that the equilibrium $\{0, 0, 0\}$ defeats all other pooling equilibria.

Proposition 4. *For sufficiently large N , the only undefeated pooling equilibrium is the equilibrium $\{0, 0, 0\}$.*

3.4 Partial-Pooling vs. Pooling

Lemma 7 highlights the fact that long-run fund performance eventually allows investors to compensate managers according to their true skill levels. That is, although managers pool in their initial choice of fund transparency, the independent information contained in their fund's return process eventually leads them to separate. This long-run outcome is of course what also happens in the partial-pooling equilibrium of section 3.2, as seen in Lemmas 4 and 5. However, there are important differences in how and how fast this long-term separation of types is achieved in the two equilibria.

Let us first consider the high types. These managers incur no transparency costs in both $\{0, \bar{t}(0), 0\}$ and $\{0, 0, 0\}$, and so only care about the speed with which investors' expectations about their fund's returns converge to μ_h . A comparison of (16) and (25) with $t = 0$ reveals that the high-type manager is better off in the partial-pooling equilibrium if $\bar{r}_n(0) > \bar{\rho}_n(0)$. This can be shown to be equivalent to

$$\lambda_h > \lambda_\ell p_G^{n-1} \frac{\mu_m}{\mu_h - \mu_m}. \quad (28)$$

Thus high types prefer the partial-pooling equilibrium when μ_h is large and p_G is small, that is, when the returns they generate are large and difficult to replicate by lower types. They also tend to prefer the partial-pooling equilibrium when μ_m is small: the addition of medium types to the pool slows down the speed at which investors conclude that the manager is a high type, and this is particularly costly when μ_m is small.

Turning our attention to the medium types, we see from (18) and (26) that they will tend to prefer the partial-pooling equilibrium if

$$\frac{\mu_m^2}{4k_t} > \frac{1}{4k_0} \left\{ [\bar{\rho}_n(0)]^2 + \left(\frac{p_G}{p_G + p_A} \right)^{n-1} \left([\bar{\rho}_n(0)]^2 - [\bar{\rho}_n(0)]^2 \right) \right\}, \quad (29)$$

where $t = \bar{t}(0)$. Since $\bar{\rho}_n(0) < \mu_m < \bar{\rho}_n(0)$ and since $k_t > k_0$, this inequality tends to hold when p_G is small, p_A is large, and k_t increases slowly in t . That is, separating from the high and low types is appealing to the medium types when it is difficult to replicate the performance of high types (i.e., p_G is small), when low types are likely to replicate the performance of medium types (i.e., p_A is large), and when transparency is not excessively costly (i.e., $\frac{dk_t}{dt}$ is small). In fact, using the refinement introduced by Mailath, Okuno-Fujiwara, and Postlewaite (1993), it is straightforward to show that the partial-pooling equilibrium $\{0, \bar{t}(0), 0\}$ defeats the pooling equilibrium $\{0, 0, 0\}$ when (29) holds.

Proposition 5. *Suppose that the inequality in (29) holds. Then the partial-pooling equilibrium $\{0, \bar{t}(0), 0\}$ defeats the pooling equilibrium $\{0, 0, 0\}$.*

4 Empirical Predictions

We now turn to the empirical implications of the model. Our predictions focus on the partial-pooling equilibrium of section 3.2. As discussed in the introduction, transparent funds naturally correspond to mutual funds while opaque funds can be thought of as hedge funds. Even within the universe of hedge funds, however, some funds reveal more information about their strategies to potential investors than others. So although our empirical predictions concentrate on contrasting mutual funds and hedge funds, one can also interpret these predictions as establishing a contrast between low- and high-transparency hedge funds.

As mentioned in section 2, μ_τ refers to the average excess return that a manager of type τ generates in a given period. If medium types separate from other types by opening (transparent) mutual funds as opposed to (opaque) hedge funds, we should expect the cross-sectional variation in the excess return of mutual funds to be lower than that of hedge funds, which are started by a pool comprised of high-skill and low-skill managers. Specifically, in our model, the alpha of a mutual fund of age n is $\alpha_n^{\text{MF}} \equiv \text{E}[\tilde{r}_n(m)] = \mu_m$ with probability one, whereas that of a hedge fund of age n is

$$\alpha_n^{\text{HF}} = \begin{cases} \text{E}[\tilde{r}_n(h) | \tilde{r}^{n-1} = r_G^{n-1}] = \mu_h > 0, & \text{prob. } \phi_n(0) \\ \text{E}[\tilde{r}_n(\ell) | \tilde{r}^{n-1} = r_G^{n-1}] = \mu_\ell = 0, & \text{prob. } 1 - \phi_n(0). \end{cases} \quad (30)$$

So clearly we have $\text{Var}(\alpha_n^{\text{HF}}) = \phi_n(0)[1 - \phi_n(0)]\mu_h > 0 = \text{Var}(\alpha_n^{\text{MF}})$. Also, since $\phi_n(0)$ converges to one over time, we expect this difference to get attenuated as funds grow older. This leads to our first prediction.

Prediction 1. *When evaluating the performance of hedge funds and mutual funds, we expect the alphas of hedge funds to be more cross-sectionally dispersed than the alphas of mutual funds. We also expect this difference in dispersion to be smaller for older (hedge and mutual) funds.*

To our knowledge, this contrast in performance heterogeneity between mutual funds and hedge funds has never been directly investigated. Agarwal, Boyson and Naik (2009) provide some evidence that on average mutual funds that employ hedge-fund-like strategies do not perform as well as the hedge funds that use similar strategies. However, they do not provide any evidence about dispersion. Another piece of evidence that appears consistent with this prediction is contained in the work by Griffin and Xu (2009), who show that hedge funds, when compared to mutual funds, seem to have wider positive and negative tails in terms of their ability to pick stocks. Finally, Sun, Wang and Zheng's (2012) findings indicate that skilled hedge fund managers use strategies that are more distinctive and profitable, while unskilled hedge fund managers tend to use more run-of-the-mill strategies that are less profitable.

For the same reason that hedge fund alphas are expected to be more cross-sectionally dispersed than mutual fund alphas, we should also expect hedge funds to fail and close more often than mutual funds. This is clearly consistent with the evidence documented by Malkiel and Saha (2005) that the attrition rate for hedge funds is three to four times that for mutual funds. Going back to the partial-pooling equilibrium of our model, we see that a fraction $1 - \phi_n(0)$ of hedge funds close in year n , whereas mutual funds always survive. Since $1 - \phi_n(0)$ is decreasing in n , we have the following prediction.

Prediction 2. *We expect the attrition rate of hedge funds to be greater than that of mutual funds of the same vintage, especially in the early years.*

In our model, as in that by Berk and Green (2004), the performance-related changes in the size of a fund originate from the updating process of investors about the manager's skill and future return prospects over time. Although the manager's compensation adjusts in a way that allows him to capture the economic rents that his skill creates, it is still the case that his expected performance gross of fees and compensation will change over time. According to our model, we should expect this adjustment to be slower for mutual funds than for hedge funds.

Prediction 3. *The average persistence in alphas is expected to be stronger in mutual funds than in hedge funds. The difference between the alphas of surviving hedge funds and mutual funds of the same vintage is expected to increase over time.*

As derived in Lemma 2, a fund's assets under management in period n are given by $A_n^{\text{HF}} = \frac{1}{2k_0} \mathbb{E}[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n]$. For a mutual fund, this quantity is constant at $A_n^{\text{MF}} = \frac{\mu m}{2k_t}$, as investors know that only medium types open mutual funds. Thus, regardless of performance in period n , the change in assets from period n to period $n + 1$ is zero for mutual funds. For hedge funds, on the other hand, the performance of the fund dictates its size in the subsequent period. Specifically, using the results of Lemma 4, we have

$$A_{n+1}^{\text{HF}} = \begin{cases} \frac{\bar{r}_{n+1}(0)}{2k_0}, & \text{if } \tilde{r}_n(\tilde{\tau}) = r_G \\ 0, & \text{otherwise,} \end{cases} \quad (31)$$

and so

$$A_{n+1}^{\text{HF}} - A_n^{\text{HF}} = \begin{cases} \frac{\bar{r}_{n+1}(0) - \bar{r}_n(0)}{2k_0} > 0, & \text{if } \tilde{r}_n(\tilde{\tau}) = r_G \\ -\frac{\bar{r}_n(0)}{2k_0} < 0, & \text{otherwise.} \end{cases} \quad (32)$$

This leads to the following empirical prediction.

Prediction 4. *We expect the relationship between performance and subsequent fund flows to be stronger in hedge funds than in mutual funds.*

Lagged fund performance and fund flows have been documented to be positively associated for mutual funds (e.g., Sirri and Tufano, 1998; Bollen and Busse, 2005) and for hedge funds (e.g., Agarwal, Daniel and Naik, 2004). To our knowledge, however, no one has contrasted the relationship across the two investment vehicles.

As mentioned above, the size of a mutual fund remains the same at $A_n^{\text{MF}} = \frac{\mu_m}{2k_t}$ in every period in our model. Thus the average size across mutual funds of the same vintage is also $\bar{A}_n^{\text{MF}} = \frac{\mu_m}{2k_t}$. However, because only well-performing hedge funds survive and because these tend to be managed by skilled managers, the average size of *surviving* funds that opened at around the same time increases over time: $\bar{A}_1^{\text{HF}} = \frac{\bar{r}_1(0)}{2k_0} < \frac{\bar{r}_n(0)}{2k_0} = \bar{A}_n^{\text{HF}}$. Thus $\bar{A}_1^{\text{HF}} - \bar{A}_1^{\text{MF}} < \bar{A}_n^{\text{HF}} - \bar{A}_n^{\text{MF}}$.

Prediction 5. *The disparity in size between hedge funds and mutual funds is expected to be more pronounced with fund age and manager tenure.*

Since compensation goes hand in hand with fund size in our model (see Lemma 2), we should expect a similar trend in per-dollar and in total manager compensation over time.

Prediction 6. *The disparity in per-dollar compensation and in total compensation between hedge fund managers and mutual fund managers is expected to increase with tenure.*

As discussed in section 3.4, the medium types are more inclined to separate from the high and low types when p_G is small. This is also precisely when the medium types choose high levels of transparency. Intuitively, when p_G is large, the low types can mimic the high types more successfully. So, to separate, the medium types do not need to impose as much of a monitoring threat on the low types. This can be seen more precisely from the condition that t must satisfy for $\{0, t, 0\}$ to be an equilibrium in our model. Specifically, recall from (17) and (19) that t must satisfy

$$(1-t) \frac{N\mu_m^2}{4k_t} = \frac{1}{4k_0} \sum_{n=1}^N p_G^{n-1} [\bar{r}_n(0)]^2. \quad (33)$$

Since the left-hand side of this condition is strictly decreasing in t but not affected by p_G and the right-hand side is strictly increasing in p_G but not affected by t , the equilibrium transparency t is decreasing in p_G . This leads to our last prediction.

Prediction 7. *The cross-sectional variation in transparency across funds is expected to be larger when fund performance is a precise estimate of skill.*

5 Conclusion

In this paper, we explore the possibility that money managers with heterogenous skills use different vehicles for the investments they make on behalf of investors. Specifically, we assume that, at the

inception of their fund, managers can choose and publicly announce the level of transparency that will apply to their fund as long it remains viable. For skilled managers, transparency comes with the benefit that investors can more quickly identify and fire their low-skill counterparts. At the same time, transparency increases the costs of managing the fund and thereby reduces the net gains that investors can expect from the fund.

In the partial-pooling equilibrium of the model, high-skill and low-skill managers pool in opaque funds, while medium-skill managers separate in transparent funds. High-skill managers tend to favor long-run performance over costly monitoring in their effort to efficiently convince investors of their quality over time. In contrast, medium-skill managers, whose expected performance may not quickly differentiate them from low-skill managers, rely on the transparency of their fund to convince investors that their skills, albeit smaller than that of high types, are still valuable. Low-skill managers, afraid that the monitoring associated with transparent funds will identify them as quacks, take their chances with opaque funds, hoping for lucky returns that will make them appear skilled. This endogenous separation of extreme and average types across funds of different transparencies produces several empirical implications that contrast the time-series and cross-sectional properties of mutual funds and hedge funds.

In his survey of the literature on hedge funds, Stulz (2007) asks (p. 176): “Since hedge funds and mutual funds essentially perform the same economic function, why do they coexist?” While we agree with Stulz that access to a wider set of financial instruments and to more complex trading strategies make hedge funds appealing to a certain class of investors,¹¹ our model leads us to question the wisdom of the recent pressure to more systematically regulate these investment vehicles. Specifically, if the opaqueness of hedge funds is a key ingredient in the efficient discovery of talent in the money management industry, then the transparency convergence prompted by the regulation of hedge funds may do more harm than good. At worst, talented money managers will even prefer working in other industries where they can more readily signal their skills and get properly rewarded for them.

¹¹In fact, the non-use of these instruments and strategies by mutual funds is partly what facilitates their monitoring and at the same time effectively increases their costs by limiting their upside potential. For more on the double-edge sword that is associated with the imposition of investment constraints on fund managers, see Almazan et al. (2004) for mutual funds, and Agarwal, Daniel and Naik (2009) for hedge funds.

Appendix

Proof of Lemma 1. The amount A_n in (6) that investors allocate to the fund follows immediately from the zero-profit condition $E[\tilde{\pi}_n | \mathcal{I}_n] = 0$. Of course, if this condition is not satisfied for any $A_n > 0$, investors prefer not to invest any money in the fund. ■

Proof of Lemma 2. The manager's compensation w_n solves the problem: $\max_{w_n} w_n A_n$, where A_n is given by (6). If $E[\tilde{r}_n(\tilde{\tau}) | \mathcal{I}_n] > 0$, the unique solution to this quadratic optimization problem is given by (8). The fund size in (9) and the manager's compensation in (10) follow then immediately from the investors' choice of A_n in (6). ■

Proof of Proposition 1. First, note that low-type managers earn zero compensation if they are identified as such by investors. The reason is that $E[\tilde{r}_n(\tilde{\tau}) | \tilde{\tau} = \ell] = \mu_\ell = 0$ for all $n \in \{1, \dots, N\}$. Further, it follows from Lemma 2 that $u_\ell(t_\ell) > 0$ if investors believe that there is a chance that the manager is of type $\tilde{\tau} = h$ or $\tilde{\tau} = m$. Thus, low-type managers strictly prefer to choose a transparency level that is also chosen by some other type. To prove the result, we are therefore left to show that $t_m = t_h$ if $t_\ell = t_m$.

We prove this result by contradiction. Suppose that $t_\ell = t_m \neq t_h$. Since $\mu_h > \mu_m > \mu_\ell$, we have $u_h(t_h) > u_h(t_m)$. Further, since high-type managers can never do worse than medium-type managers if they choose the same transparency level, we have $u_h(t_m) \geq u_m(t_m)$. But since investors correctly infer from a transparency level t_h that the manager is a high type, the manager's compensation is independent of his fund's realized returns.¹² Thus, $u_m(t_h) = u_h(t_h)$, implying that medium-type managers are strictly better off choosing a transparency level $t_m = t_h$, which contradicts our assumption that $t_m \neq t_h$. ■

Proof of Lemma 3. The investors' beliefs at the beginning of period n follow immediately from Bayes' rule. ■

Proof of Lemma 4. The expected return of a fund with a return sequence \tilde{r}^{n-1} follows immediately from the investors' beliefs specified in Lemma 3.

The expected N -period compensation of a high-type manager $u_h(t)$ in (16) follows from the period- n compensation u_n in (10) and the fact that a high-type manager generates a return of r_G in each period.

¹²As pointed out in the text, this argument appears to be inconsistent with the fact that the support of the return distribution of high types differs from that of medium types. Introducing an arbitrarily small probability that high types generate a return of r_A , however, completely resolves this apparent inconsistency.

A low-type manager, on the other hand, is identified as such if $\tilde{i}_t = 0$, which happens with probability t , and generates a return of r_G only with probability p_G . Thus, he receives the period- n compensation u_n only with probability $(1-t)p_G^{n-1}$, which implies that his expected N -period compensation $u_\ell(t)$ is given by (17). ■

Proof of Lemma 5. By choosing a transparency level $t' \neq t$, a medium-type manager identifies himself as such. From the investors' perspective, the expected return of his fund is thus μ_m in each period. His expected N -period compensation $u_m(t')$ follows then immediately from the period- n compensation u_n in (10). ■

Proof of Proposition 2. For the strategies $\mathcal{S} = \{t, t', t\}$ to be an equilibrium, the following incentive-compatibility constraints must be satisfied:

$$u_\ell(t) \geq \hat{u}_\ell(t'), \quad (\text{A1})$$

$$u_m(t') \geq \hat{u}_m(t), \quad (\text{A2})$$

$$u_h(t) \geq \hat{u}_h(t'), \quad (\text{A3})$$

where the managers' equilibrium payoffs u_τ and off-equilibrium payoffs \hat{u}_τ are defined by equations (16) to (21). Further, the managers' equilibrium payoffs also have to (weakly) exceed their payoffs from choosing any out-of-equilibrium transparency level $\hat{t} \notin \{t, t'\}$. These latter conditions are trivially satisfied for any out-of-equilibrium beliefs of investors that assign a sufficiently high probability to the event that the manager is a low type if a transparency level \hat{t} is observed.

Substituting the expressions for $u_\ell(t)$ and $\hat{u}_\ell(t')$ into (A1), the incentive-compatibility constraint for low types becomes

$$\frac{1-t}{4k_t} \sum_{n=1}^N p_G^{n-1} [\bar{r}_n(t)]^2 \geq (1-t') \frac{N\mu_m^2}{4k_{t'}}. \quad (\text{A4})$$

Since the right-hand side of (A4) is strictly decreasing in t' and is equal to zero at $t' = 1$, there always exists a $\bar{t}(t) \in (0, 1)$ such that the inequality holds for all $t' \in [\bar{t}(t), 1]$.

The incentive-compatibility constraint for medium types in (A2) can be written as

$$\frac{N\mu_m^2}{4k_{t'}} \geq \frac{1}{4k_t} \sum_{n=1}^N \left(\frac{p_G}{p_G + p_A} \right)^{n-1} [\bar{r}_n(t)]^2. \quad (\text{A5})$$

This inequality holds for N large enough since the left-hand side grows without bounds as $N \rightarrow \infty$, whereas the right-hand side is bounded above by

$$\frac{1}{4k_t} \left(1 + \frac{p_G}{p_A} \right) \mu_h^2. \quad (\text{A6})$$

Finally, the incentive-compatibility constraint for high types in (A3) can be expressed as

$$\frac{1}{4k_t} \sum_{n=1}^N [\bar{r}_n(t)]^2 \geq \frac{N\mu_m^2}{4k_{t'}}. \quad (\text{A7})$$

To see that this inequality holds for N large enough, note that the expected return $\bar{r}_n(t)$ defined in (14) converges to μ_h as $n \rightarrow \infty$. Thus, there exists an \bar{n} such that $\bar{r}_n(t) > \mu_m$ for all $n \geq \bar{n}$. This, and the fact that $\bar{r}_n(t)$ is bounded, imply that

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{4k_t} [\bar{r}_n(t)]^2 - \frac{1}{4k_{t'}} \mu_m^2 \right) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=\bar{n}}^N \left(\frac{1}{4k_t} [\bar{r}_n(t)]^2 - \frac{1}{4k_{t'}} \mu_m^2 \right) > 0 \quad (\text{A8})$$

as long as $t < t'$ and, hence, $k_t < k_{t'}$. This proves that, for N large enough, the incentive-compatibility constraint in (A3) holds for all $t < t'$. \blacksquare

Proof of Proposition 3. We first define the Mailath, Okuno-Fujiwara, and Postlewaite (1993) refinement. Consider two equilibria, $\mathcal{S} = \{t_h, t_m, t_\ell\}$ and $\mathcal{S}' = \{t'_h, t'_m, t'_\ell\}$, and let $\beta(\tau|t)$ and $\beta'(\tau|t)$ denote the investors' updated beliefs conditional on observing a transparency level t in equilibrium \mathcal{S} and \mathcal{S}' , respectively. Further, let $u_\tau(\mathcal{S})$ denote the expected payoff of the type- τ manager in equilibrium \mathcal{S} . The equilibrium \mathcal{S} *defeats* the equilibrium \mathcal{S}' if there exists a transparency level t such that the following conditions are satisfied:

- (i) There is no τ for which $t'_\tau = t$, and the set $K \equiv \{\tau | t_\tau = t\}$ is non-empty.
- (ii) For all $\tau \in K$, $u_\tau(\mathcal{S}) \geq u_\tau(\mathcal{S}')$, and for some $\tau \in K$, $u_\tau(\mathcal{S}) > u_\tau(\mathcal{S}')$.
- (iii) For some $\tau \in K$,

$$\beta'(\tau|t) \neq \frac{\lambda_\tau \pi_\tau}{\sum_{\tau' \in \{h,m,\ell\}} \lambda_{\tau'} \pi_{\tau'}}$$

for any $0 \leq \pi_\tau \leq 1$ satisfying (1) $\tau \in K$ and $u_\tau(\mathcal{S}) > u_\tau(\mathcal{S}')$ implies $\pi_\tau = 1$; and (2) $\tau \notin K$ implies $\pi_\tau = 0$.

We first show that no partial-pooling equilibrium defeats $\mathcal{S} = \{0, \bar{t}(0), 0\}$. This follows from the fact that, for N sufficiently large, the expected payoff of all three types of managers is strictly decreasing in the transparency level. For the medium type, this follows immediately from equation (18). For the high type, equation (16) implies that

$$\frac{du_h}{dt} = -\frac{dk_t/dt}{4k_t^2} \sum_{n=1}^N [\bar{r}_n(t)]^2 + \frac{1}{4k_t} \sum_{n=1}^N \frac{d[\bar{r}_n(t)]^2}{dt}. \quad (\text{A9})$$

Since k_t is strictly increasing in t and since $\bar{r}_n(t)$ converges to $\mu_h > 0$ for large n , the first term on the right-hand side of equation (A9) converges to $-\infty$ as $N \rightarrow \infty$. The second term, on the other

hand, is bounded above by $\mu_h^2 \lambda_\ell / (2k_t(1-p_G)\lambda_h)$ since $d\bar{r}_n^2/dt \leq 2\mu_h^2 p_G^{n-1} \lambda_\ell / \lambda_h$. Thus, u_h decreases in t for large N . For the low type, equation (17) lets us rewrite the manager's compensation as $u_\ell(t) = \sum_{n=1}^N p_G^{n-1} (1-t)\bar{r}_n(t)A_n/4$. Equation (14) immediately implies that $(1-t)\bar{r}_n(t)$ is decreasing in t for all $t \in [0, 1]$ and $n \in \{1, \dots, N\}$. Further, Assumption 1 implies that

$$\frac{dA_n}{dt} = \frac{\bar{r}_n(t)}{k_t^2} \left(\frac{\lambda_\ell p_G^{n-1}}{\lambda_h + \lambda_\ell(1-t)p_G^{n-1}} k_t - \frac{dk_t}{dt} \right) \leq 0, \quad \text{for all } t \in [0, 1], n \in \{1, \dots, N\}. \quad (\text{A10})$$

Thus, u_ℓ decreases in t as well. Note that this implies that $\bar{t}(t)$, which is the transparency level t' that makes (19) equal to (17), increases in t . Thus, there is no other equilibrium in which the medium type chooses a transparency level of less than $\bar{t}(0)$. We therefore have $u_m(\mathcal{S}) > u_m(\mathcal{S}')$ for $\mathcal{S}' \neq \mathcal{S}$. Similarly, there is no transparency level $t > 0$ that makes the high and the low type better off. Thus, condition (ii) cannot be satisfied and \mathcal{S} cannot be defeated by any other partial-pooling equilibrium \mathcal{S}' .

We next show that $\mathcal{S} = \{0, \bar{t}(0), 0\}$ defeats all other partial-pooling equilibria. Let $\mathcal{S}' = \{t', t'', t'\}$ denote an alternative equilibrium with $t' > 0$ and $t'' \geq \bar{t}(t')$, and consider the out-of-equilibrium (for \mathcal{S}') transparency level $\bar{t}(0)$. Condition (i) is satisfied and $K = \{m\}$. Since u_m strictly decreases in t and $t'' \geq \bar{t}(t') > \bar{t}(0)$, condition (ii) is satisfied. Condition (iii) requires the beliefs β' to satisfy $\beta'(m|\bar{t}(0)) < 1$. This must be the case in equilibrium \mathcal{S}' , since otherwise the medium type would strictly prefer the out-of-equilibrium strategy $\bar{t}(0)$ to the equilibrium strategy t'' . Thus, \mathcal{S} defeats \mathcal{S}' . ■

Proof of Lemma 6. The period- n expected return of a fund with a return sequence \tilde{r}^{n-1} follows immediately from Bayes' rule. ■

Proof of Lemma 7. The expected N -period compensation of a high-type manager $u_h(t)$ in (25) follows from the period- n compensation u_n in (10) and is based on the expected return $\bar{\rho}_n(t)$ since a high-type manager always generates a return of r_G .

A medium-type manager receives the same period- n compensation as long as his realized return is equal to r_G , which happens with probability $p_G/(p_G + p_A)$ in each period. If investors observe a return of r_A , they know that the manager is not a high type and the manager's compensation from then on is based on the expected return $\bar{\rho}_n(t)$.

A low-type manager is identified as such if $\tilde{t}_t = 0$, which happens with probability t . He generates a return of r_G with probability p_G . Thus, his period- n compensation is based on the expected return $\bar{\rho}_n(t)$ with probability $(1-t)p_G^{n-1}$. As long as he can avoid a return of r_B , his compensation is based on the expected return $\bar{\rho}_n(t)$, which happens with probability $(1-t)(p_G + p_A)^{n-1}$ in period n . If his fund realizes a return of r_B , investors know that he is a low type and he receives zero compensation from then on. ■

Proof of Proposition 4. The definition of the Mailath, Okuno-Fujiwara, and Postlewaite (1993) refinement can be found in the proof of Proposition 3.

We first show that no pooling equilibrium defeats $\mathcal{S} = \{0, 0, 0\}$. Analogous arguments to those made in the proof of Proposition 3 show that, for large N , the expected payoff of all three types of managers is strictly decreasing in the transparency level. Thus, condition (ii) of the equilibrium refinement cannot be satisfied for any other pooling equilibrium with a transparency level of $t > 0$.

We next show that $\mathcal{S} = \{0, 0, 0\}$ defeats all other pooling equilibria. Let $\mathcal{S}' = \{t', t', t'\}$ denote an alternative equilibrium with $t' > 0$, and consider the out-of-equilibrium (for \mathcal{S}') transparency level $t = 0$. Condition (i) of the equilibrium refinement is satisfied and $K = \{h, m, \ell\}$. Since u_h , u_m , and u_ℓ strictly decrease in the transparency level, condition (ii) is satisfied. Condition (iii) requires the beliefs $\beta'(\tau|t = 0)$ to be different from the prior beliefs λ_τ for some type $\tau \in \{h, m, \ell\}$. This must be the case in equilibrium \mathcal{S}' , since otherwise a manager of type τ would strictly prefer the out-of-equilibrium strategy $t = 0$ to the equilibrium strategy t' . Thus, \mathcal{S} defeats \mathcal{S}' . ■

Proof of Proposition 5. The definition of the Mailath, Okuno-Fujiwara, and Postlewaite (1993) refinement can be found in the proof of Proposition 3.

Let \mathcal{S} denote the partial-pooling equilibrium $\{0, \bar{t}(0), 0\}$ and \mathcal{S}_p denote the pooling equilibrium $\{0, 0, 0\}$. Consider the out-of-equilibrium (for \mathcal{S}_p) transparency level $\bar{t}(0)$. Condition (i) of the equilibrium refinement is satisfied and $K = \{m\}$. If the inequality in (29) holds, $u_m(\mathcal{S}) > u_m(\mathcal{S}_p)$ and thus condition (ii) is satisfied as well. Condition (iii) requires the beliefs β_p to satisfy $\beta_p(m|\bar{t}(0)) < 1$. This must be the case in equilibrium \mathcal{S}_p , since otherwise the medium type would strictly prefer the out-of-equilibrium strategy $\bar{t}(0)$ to the equilibrium strategy $t = 0$. Thus, \mathcal{S} defeats \mathcal{S}_p . ■

References

- Agarwal, Vikas, Nicole M. Boyson, and Narayan Y. Naik, 2009, "Hedge Funds for Retail Investors? An Examination of Hedged Mutual Funds," *Journal of Financial and Quantitative Analysis*, 44(2), 273-305.
- Agarwal, Vikas, Naveen D. Daniel, Narayan Y. Naik, 2004, "Flows, Performance, and Managerial Incentives in Hedge Funds," Working Paper, Georgia State University.
- Agarwal, Vikas, Naveen D. Daniel, Narayan Y. Naik, 2009, "Role of Managerial Incentives and Discretion in Hedge Fund Performance," *Journal of Finance*, 64(5), 2221-2256.
- Aggarwal, Rajesh K., and Philippe Jorion, 2011, "Is There a Cost to Transparency?" Working Paper, University of Minnesota.
- Almazan, Andres, Keith C. Brown, Murray Carlson, and David A. Chapman, 2004, "Why Constrain Your Mutual Fund Manager?" *Journal of Financial Economics*, 73(2), 289-321.
- Alós-Ferrer, Carlos, and Julien Prat, 2011, "Job Market Signaling and Employer Learning," *Journal of Economic Theory*, forthcoming.
- Berk, Jonathan B., and Richard C. Green, 2004, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269-1295.
- Bhattacharya, Sudipto, and Paul Pfleiderer, 1985, "Delegated Portfolio Management," *Journal of Economic Theory*, 36(1), 1-25.
- Bollen, Nicolas P. B., and Jeffrey A. Busse, 2005, "Short-Term Persistence in Mutual Fund Performance," *Review of Financial Studies*, 18(2), 569-597.
- Bookstaber, Richard, 2003, "Hedge Fund Existential," *Financial Analysts Journal*, 59(5), 19-23.
- Brown, Stephen J., and William N. Goetzmann, 2003, "Hedge Funds with Style," *Journal of Portfolio Management*, 29(2), 101-112.
- Brown, Stephen J., William N. Goetzmann, Bing Liang, and Christopher Schwarz, 2008, "Mandatory Disclosure and Operational Risk: Evidence from Hedge Fund Registration," *Journal of Finance*, 63(6), 2785-2815.

Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey D. Kubik, 2004, "Does Fund Size Erode Mutual Fund Performance? The Role of Liquidity and Organization," *American Economic Review*, 94(5), 1276-1302.

Cho, In-Koo, and David M. Kreps, 1987, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179-221.

Daley, Brendan, and Brett Green, 2011, "Market Signaling with Grades," Working Paper, Duke University.

Dangl, Thomas, Yuchang Wu, and Josef Zechner, 2008, "Market Discipline and Internal Governance in the Mutual Fund Industry," *Review of Financial Studies*, 21(5), 2307-2343.

Das, Sanjiv Ranjan, and Rangarajan K. Sundaram, 2002, "Fee Speech: Signaling, Risk-Sharing, and the Impact of Fee Structures on Investor Welfare," *Review of Financial Studies*, 15(5), 1465-1497.

Feltovich, Nick, Richmond Harbaugh, and Ted To, 2002, "Too Cool for School? Signalling and Countersignalling," *RAND Journal of Economics*, 33(4), 630-649.

Fishman, Michael J., and Kathleen M. Hagerty, 2003, "Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers," *Journal of Law, Economics, and Organization*, 19(1), 45-63.

Frank, Mary Margaret, James M. Poterba, Douglas A. Shackelford, and John B. Shoven, 2004, "Copycat Funds: Information Disclosure Regulation and the Returns to Active Management in the Mutual Fund Industry," *Journal of Law and Economics*, 47(2), 515-541.

Fung, William, and David A. Hsieh, 1997, "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds," *Review of Financial Studies*, 10(2), 275-302.

Gervais, Simon, Anthony W. Lynch, and David K. Musto, 2005, "Fund Families as Delegated Monitors of Money Managers," *Review of Financial Studies*, 18(4), 1139-1169.

Gomes, Armando, 2000, "Going Public without Governance: Managerial Reputation Effects," *Journal of Finance*, 55(2), 615-646.

Griffin, John M., and Jin Xu, 2009, "How Smart Are the Smart Guys? A Unique View from Hedge Fund Stock Holdings," *Review of Financial Studies*, 22(7), 2531-2570.

- Grossman, Sanford J., and Motty Perry, 1986, "Perfect Sequential Equilibrium," *Journal of Economic Theory*, 39(1), 97-119.
- Huberman, Gur, and Shmuel Kandel, 1993, "On the Incentives for Money Managers: A Signalling Approach," *European Economic Review*, 37(5), 1065-1081.
- Huddart, Steven, 1999, "Reputation and Performance Fee Effects on Portfolio Choice by Investment Advisers," *Journal of Financial Markets*, 2(3), 227-271.
- Koski, Jennifer Lynch, and Jeffrey Pontiff, 1999, "How Are Derivatives Used? Evidence from the Mutual Fund Industry," *Journal of Finance*, 54(2), 791-816.
- Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite, 1993, "Belief-Based Refinements in Signalling Games," *Journal of Economic Theory*, 60(2), 241-276.
- Malkiel, Burton G., and Atanu Saha, 2005, "Hedge Funds: Risk and Return," *Financial Analysts Journal*, 61(6), 80-88.
- Sirri, Erik R., and Peter Tufano, 1998, "Costly Search and Mutual Fund Flows," *Journal of Finance*, 53(5), 1589-1622.
- Spence, Michael, 1973, "Job Market Signaling," *Quarterly Journal of Economics*, 87(3), 355-374.
- Spiegel, Yossef, and Daniel F. Spulber, 1997, "Capital Structure with Countervailing Incentives," *RAND Journal of Economics*, 28(1), 1-24.
- Stein, Jeremy C., 2005, "Why Are Most Funds Open-End? Competition and the Limits of Arbitrage," *Quarterly Journal of Economics*, 120(1), 247-272.
- Stulz, René M., 2007, "Hedge Funds: Past, Present, and Future," *Journal of Economic Perspectives*, 21(2), 175-194.
- Sun, Zheng, Ashley Wang, and Lu Zheng, 2012, "The Road Less Traveled: Strategy Distinctiveness and Hedge Fund Performance," *Review of Financial Studies*, 25(1), 96-143.
- Taylor, Curtis R., 1999, "Time-on-the-Market as a Sign of Quality," *Review of Economic Studies*, 66(3), 555-578.