Trading on Short-Term Information

by

ALEXANDER GÜMBEL*

This paper shows that investors may want fund managers to acquire and trade on short-term instead of more profitable long-term information. This improves learning about managerial ability from performance observations, for two reasons. Firstly, short-term information is of higher quality, which allows the investor to draw sharper inferences over a manager's type. Secondly, performance observations under long-term informed trade are contaminated by noise contained in prices, which further weakens inference. The paper thus explicitly links the degree of short-term information dissemination to the profitability and the learning implications of short-term *versus* long-term informed trading. (JEL: D 82, D 83, G 14, G 23)

1 Introduction

There has been considerable debate, among economists and practitioners alike, concerning short-termism in financial markets. Short-termism refers to factors concerning the near future, like short-term profits and cash flows, carrying excessive weight in decision-making compared to factors regarding the longer term. Excessive is here defined relative to a first-best benchmark prevailing in a frictionless economy. Short-termism among fund managers frequently is held responsible for undervaluation of and therefore underinvestment in assets with long-horizon cash flows. In turn, short-termist economies are alleged to suffer unnecessarily low growth rates (MARSH [1990]).[1]

---

[1] DOW AND GORTON [1994] and ALLEN, MORRIS, AND SHIN [2004] provide a formal treatment of why prices may not reflect long-term information effectively when traders have short horizons. GÜMBEL [2005] explores the welfare implications of such a failure. Note also that short-termism may arise in the absence of agency problems, for example due to risk aversion (HOLDEN AND SUBRAHMANYAM [1996]).

SHLEIFER AND VISHNY [1997] argue that fund managers forgo long-term arbitrage opportunities if investors remove funds from them after poor performance. When stock prices are noisy, poor interim performance is a risk inherent in long-term arbitrage, which managers will therefore seek to avoid. An interesting question is then what the investors' optimal behaviour in such a setting should be. While a number of papers have investigated optimal-contracting issues in delegated portfolio management (see BHATTACHARYA AND PFLEIDERER [1985], STOUGHTON [1993], HEINKEL AND STOUGHTON [1994]) and how contracts affect trading strategies (*e.g.*, ALLEN AND GORTON [1993], DOW AND GORTON [1997], DASGUPTA AND PRAT [2003]), it remains an open issue how contracting between the investor and the fund manager affects the information horizon of trading strategies.

I show that an investor may prefer trade on short-term information by the manager, because this allows her to learn more quickly about the manager's unknown ability. This comes at the cost of smaller expected trading profits from acquiring short-term information. The basic idea is captured in a setting where a high-ability manager sometimes learns that an asset is over- or undervalued from a private signal indicating that a future dividend payment will be higher or lower than expected by the market. Low-type managers never spot misvaluations. The investor can choose whether she wants a manager of unknown ability to acquire information about a dividend payment that is expected in the long term (two periods ahead) or in the short term (one period ahead). Moreover, I assume that it is easier to predict the short-term future and therefore a high-type manager is more likely to identify a short-term than a long-term misvaluation.

In my model, acquiring the higher-quality short-term signal may be less profitable for two reasons. Firstly, a manager who identifies a long-term misvaluation has two trading dates available before the dividend payment actually occurs. He can therefore profit from his private information twice, compared to only once when he acquires short-term information. Secondly, I assume that short-term information is partially disseminated into the price, because there is another short-term speculator who sometimes identifies short-term misvaluations. A short-term informed manager then competes in the financial market with this speculator, which reduces the rent from trade on private information.

On the other hand, the investor learns more quickly about the manager's ability when he acquires short-term information, because it is inherently of higher quality than long-term information. High-type managers are less likely to miss an arbitrage opportunity when they try to identify short-term than long-term misvaluations. This then makes it easier for the investor to identify a high-type manager when he acquires short-term information. Performance observations generated from the acquisition of the higher-quality short-term signal essentially allow sharper inferences about the manager's ability than the inferences drawn from trade on long-term information.

Existing literature on short-termism has largely focused on the problem of investment decisions by firm managers (VON THADDEN [1995], SHLEIFER AND VISHNY [1990], STEIN [1989]). These papers show that short-termism may be the result of

the agency problem between the firm and its manager. Instead I show that short-termist behaviour may prevail in the absence of a moral-hazard problem over the choice of horizon. In a related paper VON THADDEN [1995] shows that short-term investment projects may be desirable, because they provide better interim information about project quality. This insight does not carry across to the fund-management context in a straightforward manner, because interim performance information from a long-horizon trading strategy may be available even in the short run. As DEMIRAG [1995, pp. 43f.] points out:

"It is ... reasonable to argue that pressures to maximise short-run returns ... are in principle compatible with a willingness to ignore short-term cash flows, profits, and dividends in favour of long-term prospects. A fund manager who consistently recognized such prospects and invested accordingly, shortly before others did, would 'perform' extremely well in the short-term."

I model explicitly a setup where managers who acquire long-term information may perform well in the short run, because prices sometimes reflect future misvaluations before they are realized. Nevertheless, short-term information acquisition may be preferred by the investor, because it confers a learning advantage about managerial ability on her.

The remainder of the paper is organized as follows. Section 2 lays out the basic model of privately informed trade, and section 3 provides its solution. In section 4 the investor's full optimization problem is presented, and its solution is derived in section 5. Section 6 discusses robustness in the presence of moral hazard. Section 7 concludes. The Appendix contains the proofs.

2   The Model

The economy is in discrete time $t = 0, 1, ..., \infty$. There are dividend-paying assets, and for simplicity each asset has exactly one date at which it may pay an uncertain dividend $d_t \in \{-1, 0, 1\}$ with probabilities 1/4, 1/2, and 1/4, respectively. Assume also that at each date there is exactly one asset for which that date is the payout date, *i.e.*, an asset is uniquely defined by the date t at which its dividend is due. I refer to the asset for which date t is the dividend payout date simply as asset t. There is also a risk-free asset with return r.

An infinite-lived investor can employ a fund manager to engage in active portfolio management. The manager can acquire a private signal about the future dividend payment, which is then used for trading decisions. There are two types of managers $m \in \{L, H\}$, and only the high-type manager may receive an informative signal. Neither the investor nor the managers can observe a manager's type. There are infinitely many managers, and a fraction γ_0 of them are high types $m = H$. In order to simplify the analysis of the learning problem I assume that there is no agency problem between the investor and the fund manager. Section 6 provides

a discussion of robustness when an agency problem between investor and fund manager exists.

2.1 Information Structure

At any given date t a manager can produce *either* short-term information about the dividend payment next period (d_{t+1}) *or* long-term information about the dividend payment two periods from now (d_{t+2}). He then receives a signal $s_t^{t+h_t} \in \{-1, 0, 1\}$, where $h_t = 1$ $(h_t = 2)$ indicates that short-term (long-term) information is acquired. For expositional purposes assume that the signal is always zero when no information is acquired, *i.e.*, $s_t^{t+1} = 0$ if $h_t = 2$ and $s_t^{t+2} = 0$ if $h_t = 1$.

The signal quality differs between high- and low-type managers. Assume that low types only ever receive the signal $s_t^{t+h_t} = 0$. High-type managers on the other hand receive a signal $s_t^{t+h_t} = 0$ whenever the dividend payment is equal to its mean value of zero, *i.e.*, $\text{prob}(s_t^{t+h_t} = 0 | d_{t+h_t} = 0) = 1$. When the dividend payment is nonzero, high types can learn this in advance with probability v_s for short-term and v_l for long-term information, *i.e.*,

$$v_s \equiv \text{prob}(s_t^{t+1} = 1 | d_{t+1} = 1) = \text{prob}(s_t^{t+1} = -1 | d_{t+1} = -1) \quad \text{and}$$
$$v_l \equiv \text{prob}(s_t^{t+2} = 1 | d_{t+2} = 1) = \text{prob}(s_t^{t+2} = -1 | d_{t+2} = -1),$$

respectively. Moreover, with complementary probability $1 - v_s$ (or $1 - v_l$) a high type receives the zero signal $s_t^{t+h_t} = 0$, *i.e.*, he never observes $s_t^{t+h_t} = -1$ when $d_{t+h_t} = 1$ and *vice versa*. Hence, $1 - v_s = \text{prob}(s_t^{t+1} = 0 | d_{t+1} = 1) = \text{prob}(s_t^{t+1} = 0 | d_{t+1} = -1)$, and similarly for the long-term signal. Moreover, assume that it is easier to predict the near-term future than the longer-term future, *i.e.*, $v_s > v_l$.

In line with the assumption that there are no agency problems between managers and the investor, one could think of the manager purely as a signal generator, whose output is directly observable by the investor. The investor can then use the signals to trade in the financial market.

Moreover, there is one short-term speculator for each asset. A short-term speculator for asset $t + 1$ receives private (short-term) information at the previous date. His signal $\bar{s}_t^{t+1}$ has the same structure as a high-type fund manager's, except that the probability of receiving a signal indicating a nonzero dividend is captured by the parameter $\beta \equiv \text{prob}(\bar{s}_t^{t+1} = 1 | d_{t+1} = 1) = \text{prob}(\bar{s}_t^{t+1} = -1 | d_{t+1} = -1)$.

The short-term speculator serves two purposes. Firstly, his presence renders trade on short-term information less profitable than trade on long-term information, because the investor loses her information monopoly. Secondly, prices become more informative as time goes by. In particular, they may reflect information about a dividend payment one period before it is due. The investor can therefore evaluate the manager after *one* period, even when he acquired (long-term) information regarding the dividend payment *two* periods ahead. The degree to which the price allows the investor to make an inference about the manager's ability depends on its informativeness, which is increasing in the short-term speculator's quality of information, as measured by β.

2.2 *The Financial Market*

There are four types of potential traders in each asset: The investor, a noise trader, a short-term speculator, and a market maker. Trade and price formation are modelled in a setup based on KYLE [1985]. All agents wishing to trade at a given date submit their (market) orders simultaneously to the risk-neutral market maker, who sets the price and balances order flow out of his inventory.

Define q_t^τ as the investor's order at date t for asset τ, and similarly define $\overline{q}_t^{t+1}$ as the short-term speculator's order for asset $t+1$. The noise trader's date-t demand for asset $\tau \in \{t+1, t+2\}$ is given by $n_t^\tau \in \{-1, 1\}$. Each realization of n_t^τ is assumed to be equally likely, and demand is uncorrelated across time and across assets. For $\tau > t+2$ noise-trader demand is zero. It would be fairly straightforward to endogenize noise-trader demand, for example by modelling it as resulting from a hedging need (see SPIEGEL AND SUBRAHMANYAM [1992]). In order to focus the analysis, however, noise-trader demand is treated as exogenous here.

At date t the following sequence of events occurs. First, the dividend payment d_t is made to the holders of asset t. Second, the short-term speculator for asset $t+1$ receives his private signal $\overline{s}_t^{t+1}$, and the investor simultaneously receives a short-term signal s_t^{t+1}. Orders for short-term trades q_t^{t+1}, $\overline{q}_t^{t+1}$, and n_t^{t+1} are then submitted to the market maker for asset $t+1$, who sets a price p_t^{t+1} and executes the orders. Finally, the investor receives a signal s_t^{t+2}, and orders for long-term trades q_t^{t+2} and n_t^{t+2} are submitted to the market maker for asset $t+2$, who sets a price p_t^{t+2} and executes the orders.[2] The investor and the short-term speculator can condition their date-t orders on past public information (prices and dividend payments) and the private signal(s) each received up to the date-t signals $s_t^{t+h_t}$ and $\overline{s}_t^{t+1}$, respectively. Remember that at any date t, only one signal, s_t^{t+1} or s_t^{t+2}, may ever be nonzero, depending on the information horizon h_t.

The short-term speculator active in asset $t+1$ submits his order so as to maximize the expected present value of trading profits, conditional on his signal and past prices:

$$\max_{\overline{q}_t^{t+1}} E\left[\overline{q}_t^{t+1}\left(\frac{d_{t+1}}{1+r} - p_t^{t+1}\right) \mid \overline{s}_t^{t+1}, p_{t-1}^{t+1}\right].$$

There is exactly one market maker for each asset. Suppose that the market maker can observe each individual order submitted to him, but he does not observe who submitted the order. The total order flow Q_t^τ is therefore the set of nonzero orders submitted at date t for asset τ.[3] In case of long-term orders Q_t^{t+2} there are only two possible traders: the noise trader and the investor. For short-term orders Q_t^{t+1} there is also the short-term speculator, increasing the maximum number of orders to three.

---

[2] The reason for introducing a short delay between short-term and long-term trades is to allow the investor to change a manager at date t on the basis of date-t price movements without losing out on a round of trade. This will only become relevant in section 4 when the investor's employment policy is analysed.

[3] In KYLE [1985] the market maker observes the sum of individual orders, which is more convenient when working with normal distributions. See DOW AND GORTON [1997] for a treatment similar to ours.

As in KYLE [1985], I construct an equilibrium with informationally efficient prices. By this is meant that prices are equal to the expected discounted asset value conditional on the market maker's information at the time t of setting the price, *i.e.*,

$$p_t^\tau = \frac{1}{(1+r)^{\tau-t}} E\left[d_\tau | F_t\right],$$

where F_t is the market maker's information set at date t. One way to derive the efficiency condition from first principles would be to allow for Bertrand competition between many market makers for order-flow execution. In KYLE [1985], and most of the literature building on it, this extra step is removed from the analysis for simplicity's sake. Instead, equilibrium is constructed by replacing the corresponding optimization problem of multiple market makers with the simple informational efficiency condition for the price as above. For that reason, Kyle's model is, in his own words, "not quite a game theoretic one" (p. 1318).

In order to apply standard game-theoretic equilibrium concepts, it is necessary to model the market makers as payoff-maximizing agents. I therefore choose a different shortcut that achieves informationally efficient prices with fully endogenous market maker behaviour, without having to add a full analysis of market maker competition. In particular, assume that each market maker has a payoff function that is strictly decreasing in his pricing error. This rewards him directly for setting informationally efficient prices. The payoff function of the market maker for asset τ is given by

$$M^\tau = -\sum_{t'=1}^{2} \left(\frac{d_\tau}{(1+r)^{t'}} - p_{\tau-t'}^\tau\right)^2.$$

Note that the solution to $\max_{p_{\tau-t'}^\tau} E\left[M^\tau | F_{\tau-t'}\right]$ is given by

$$p_{\tau-t'}^\tau = E\left[\frac{d_\tau}{(1+r)^{t'}} | F_{\tau-t'}\right],$$

which yields informationally efficient prices.[4]

Finally, assume that no side payments between the investor, the market maker, and the short-term trader are possible. This is to rule out that the investor tries to collude with the market maker and bribe him to set prices inefficiently at the expense of the noise traders. Similarly, the assumption rules out that the investor colludes with the short-term speculator to act as an information monopolist, again at the expense of the noise traders.

3 Trading Strategies and Profits

This section considers the Perfect Bayesian Equilibrium (PBE) of the restricted trading game described above. In this restricted game the investor's only choice variables at each date t are the orders q_t^{t+1} and q_t^{t+2}. She chooses orders so as to maximize the expected present value of trading profits conditional on her signals

---

[4] Note that this pricing strategy yields zero expected profits.

s_t^{t+1}, s_t^{t+2}, and s_{t-1}^{t+1} (signals about assets before $t + 1$ are irrelevant at date t, because the dividend payment has already been made and no further trade occurs). Hence, the restricted game examined in this section differs from the full game described in section 4 by allowing for neither managerial hiring and firing decisions, nor the choice of an information horizon.

The analysis of the restricted game serves to simplify the analysis of the full game of sections 4 and 5. It also provides a comparison of trading profits from short-term and long-term information acquisition. This yields a parameter restriction on when long-term information acquisition is the first-best choice of information horizon. This is obviously important, given that the main result of this paper will be to show that short-term information acquisition may be chosen even when long-term information acquisition is the first best.

Assume that it is common knowledge that $\gamma_t > 0$.

PROPOSITION 1 *There exists a PBE of the trading game. The investor's equilibrium trading strategy for all $t \geq 0$ is given by*

(1) $$q_t^{t+2} = s_t^{t+2},$$

(2) $$q_t^{t+1} = \begin{cases} s_{t-1}^{t+1} & \text{if } s_{t-1}^{t+1} \neq 0, \ t > 0, \\ s_t^{t+1} & \text{otherwise,} \end{cases}$$

and the equilibrium strategy of the short-term speculator active in asset $t + 1$ is given by

(3) $$\overline{q}_t^{t+1} = \overline{s}_t^{t+1}.$$

The investor's trading strategy is straightforward: she buys after positive news, sells after negative news, and does not trade after the zero signal. Moreover, if the investor has received an informative long-term signal $s_{t-1}^{t+1} \neq 0$, then at t she trades again in the same direction as before. If she has not received an informative long-term signal at $t - 1$, she follows the date-t short-term signal. Similarly, the short-term speculator buys (sells) after positive (negative) news and does not trade after the zero signal.

In the equilibrium of Proposition 1, the investor and the short-term speculator sometimes trade at date t in asset $t + 1$ even though the date-$t - 1$ price fully revealed private information ($Q_{t-1}^{t+1} \in \{\{-1, -1\}, \{1, 1\}\}$). Such trades yield zero profits. There exist other equilibria in which the investor and/or the short-term speculators do not trade in this case. Note that both equilibrium prices and trading profits are unchanged in these alternative equilibria. Outside of these trivially different equilibria, there is no other equilibrium (see the proof of Proposition 2 for a formal treatment of this issue).

Consider an investor who employs a fund manager who is a high type with probability γ_t. We now calculate the total trading profits from (potentially repeated) trade in one asset as a function of the information horizon h_t. In the case of long-term information ($h_t = 2$) this means calculating the expected discounted profits from trading q_t^{t+2} and q_{t+1}^{t+2} according to the equilibrium in Proposition 1. In the case

of short-term information ($h_t = 1$), profits accrue from the trade q_t^{t+1} alone. For short-term information acquisition consider the case where information about asset $t + 1$ is acquired for the first time at date t. This may be different from repeated acquisition of information, because expected trading profits would otherwise depend on the previous date's signal realization. Although it is straightforward to calculate expected trading profits even then, it turns out not to be necessary: it is shown below that acquiring short-term information after long-term information is never optimal.

PROPOSITION 2 *The date-t present value of the investor's expected trading profits from long-term information acquisition at date t in any PBE of the trading game is given by*

$$(4) \qquad E\left[\pi\left(\gamma_t, h_t = 2\right)\right] = \frac{\gamma_t v_l}{4\left(1 + r\right)^2}\left(1 + \frac{1 - \beta}{2}\right).$$

The date-t present value of the investor's expected trading profits from first time short-term information acquisition at date t in any PBE of the trading game is given by

$$(5) \qquad E\left[\pi\left(\gamma_t, h_t = 1\right)\right] = \frac{\gamma_t v_s\left(1 - \beta\right)}{4\left(1 + r\right)}.$$

Even though there are multiple equilibria of the trading game, profits are uniquely determined by (4) and (5). This is because the equilibria differ in a trivial way: at some nodes players are indifferent between actions that yield zero payoffs. Hence, trading profits are not affected across the possible equilibria.

Expected trading profits increase with the manager's reputation γ_t and with the quality v_l or v_s of a high-type manager's signal, and they decrease with increasing discount rate r and with increasing quality β of the short-term speculator's information. Acquiring the long-term signal and trading on it has two advantages over using the short-term signal. Firstly, when the investor first trades on long-term information, she is the only trader endowed with this information. When she trades on the short-term signal she competes with another potentially informed speculator, which reduces her rent from trade on this information. Secondly, the investor has the opportunity to trade on long-term information twice, viz., if the mispricing persists after her first trade, she has a further arbitrage opportunity in the following period. The expected profits in the second round of trade then are decreasing in β. However, since β does not affect the profits of the first round of trade, profits from long-term information acquisition are less sensitive to a change in β than those from short-term information acquisition. Thus, there exists a value of β high enough that the expected trading profits from acquiring long-term information exceed those from short-term information. A simple comparison between (4) and (5) yields the following result:

COROLLARY 1 *Acquiring long-term information is more profitable than acquiring short-term information if*

$$(6) \qquad \beta \geq \frac{v_s\left(1 + r\right) - \frac{3}{2}v_l}{v_s\left(1 + r\right) - \frac{1}{2}v_l}.$$

During the remainder of the paper I focus on parameter values satisfying (6) and $v_s > v_l$. Hence, I focus on a setting where short-term information is of higher quality than long-term information, but trading profits from short-term information are lower, because another speculator also trades on short-term information. The more informed this other speculator is (the higher β), the less advantageous it is to acquire short-term compared to long-term information.

Note also, that the condition (6) is independent of a manager's reputation γ_t. By assuming that (6) holds, it is therefore also assumed that long-term information acquisition is optimal in the benchmark case where the manager's type is known to the investor.

We can apply the results of Proposition 2 and write the expected discounted trading profits from one-off information acquisition h_t as a function of the manager's reputation and the chosen information horizon as follows:

$$
(7) \qquad E[\pi(\widehat{\gamma}_t, h_t)] = \begin{cases} \dfrac{\widehat{\gamma}_t v_s (1-\beta)}{4(1+r)} & \text{if } h_t = 1, \\[2ex] \dfrac{\widehat{\gamma}_t v_l}{4(1+r)^2}\left(1 + \dfrac{1-\beta}{2}\right) & \text{if } h_t = 2. \end{cases}
$$

4 Learning from Performance Observations

This section completes the description of the full game. In addition to the trading decision, the investor now chooses an employment policy and an information horizon at each date. These decisions affect contemporaneous profits as well as the investor's ability to learn the manager's type. Moreover, we need to consider more carefully the investor's payoffs, which in general are affected by trading profits and the wage paid to the fund manager.

First, consider wage determination for the fund manager. If a manager is picked from the pool, he has no bargaining power and receives the minimum wage of zero from the investor. If a manager has been employed for at least one period and established a better reputation, he may gain some bargaining power. Under the assumption that a manager's employment record is publicly observable and that there are other investors who could hire the manager, he could extract the entire surplus from an increased reputation by threatening to switch to a different investor. It is assumed, however, that the manager has an element of investor-specific ability that is nontransferable to other investors. A manager, for example, who has proven himself at managing a portfolio of UK listed shares is not necessarily equally able to manage an East Asia hedge fund. A fraction of the informational rent can therefore be captured by the investor.

Assume therefore that a higher-reputation manager can capture a fraction α of his trading profit. Assume that α is such that the manager and the investor are both *strictly* better off from dealing with each other than from ending the relationship. This assumption ensures that the investor has an interest in identifying a high-type manager. Otherwise the investor would be indifferent between employing either type of manager, and the employment policy would be irrelevant. Empirical evidence,

however, suggests that investors react to past managerial performance by switching between funds or laying off poorly performing managers (see CHEVALIER AND ELLISON [1997], [1999]).

We can then write the expected discounted payoff to the investor from acquiring information of horizon h_t and trading according to (1) and (2), using (7), as

$$(8) \qquad E[\Pi(\widehat{\gamma}_t, h_t)] = \begin{cases} E[\pi(h_t, \widehat{\gamma}_t)] & \text{if } \widehat{\gamma}_t \leq \gamma_0, \\ (1 - \alpha)\, E[\pi(h_t, \widehat{\gamma}_t)] & \text{if } \widehat{\gamma}_t > \gamma_0. \end{cases}$$

At date t the investor chooses an employment variable $e_t \in \{0, 1\}$. She can choose either to fire the incumbent manager and employ a new manager from the pool ($e_t = 0$), or to retain the incumbent manager ($e_t = 1$). Simultaneously the investor chooses an information horizon h_t. She can do so after having observed the dividend d_t, or she can wait until after having observed the price p_t^{t+1}. This allows an investor to observe the date-t performance of a manager who acquired long-term information at date $t-1$, by observing a price p_t^{t+1} that may reflect information on d_{t+1}. The investor can then fire the manager if she wishes to and hire a new manager. Since by assumption (see section 3) the market for the long-term asset $t+2$ only meets with a short delay, the new manager can immediately generate a further long-term signal on which the investor can trade at the same date t.[5]

The investor can then update her belief about the manager's type by comparing the manager's date-$t-1$ signal $s_{t-1}^{t-1+h_{t-1}}$ with the dividend realization and/or the realized date-t price. Denote by γ_t the updated probability that the manager, employed between dates $t-1$ and t, is a high type, conditional on the date-t performance observations. The relationship between reputation updates and performance observations is fully developed in sections 4 and 5. Moreover, denote by $\widehat{\gamma}_t$ the reputation of the manager who is acquiring information at date t. If the incumbent manager is retained, we have that $\widehat{\gamma}_t = \gamma_t$. If he is fired and replaced by a new manager, we have that $\widehat{\gamma}_t = \gamma_0$. To summarize:

$$(9) \qquad \widehat{\gamma}_t\,(\gamma_t, e_t) = \begin{cases} \gamma_t & \text{if } e_t = 1, \\ \gamma_0 & \text{if } e_t = 0. \end{cases}$$

4.1 The Optimal Policy

The investor faces a dynamic optimization problem with choice variables e_t, h_t and q_t^{t+1}, q_t^{t+2} at each date t. The value function is the expected present discounted value of all future payoffs. An equilibrium of the full game is now defined as a PBE such that at each date, the investor chooses actions e_t, h_t and q_t^{t+1}, q_t^{t+2} so as to maximize her value function. Payoffs and choice variables of all other players remain unchanged.

---

[5] This assumption clearly increases the value of long-term information acquisition due to learning about managerial ability. It therefore makes it harder to show the main result that short-term may be preferred to long-term information acquisition.

LEMMA 1 *The equilibrium trading and price-setting strategies of the full game are the same as those of the restricted trading game.*

Lemma 1 says that the trading and price-setting strategies of the overall equilibrium are simply given by the equilibrium strategies of the trading game.[6] The intuition for this is simple: the investor can separate the trading decision from the employment and information-horizon decisions. She therefore continues to choose trades so as to maximize expected trading profits contingent on the signal realizations. This allows us to write the investor's overall optimization problem in a much simpler form. In particular, at each date we can use the expected payoff from choosing a particular trading horizon as the instantaneous payoff at that date. Moreover, we can ignore the choice of trades at each date, since we know how they are determined. We can then reduce the investor's optimization problem to a choice over information horizon and employment policy at each date, where the instantaneous payoff from this choice is given by the expected discounted payoff of the reduced-form game of trade in the corresponding asset.

The investor's immediate payoff in each of the trading games is uniquely given by (8). The overall game can therefore be solved by solving the reduced-form game in which a particular information horizon and employment variable directly generate an expected payoff given by (8). This game features no further strategic interaction and is a pure decision problem in which the investor chooses an information horizon and an employment variable at each date.

The investor's value function is the sum of all the expected payoffs given by (8). This can be written as a function of two state variables.[7] The first is the incumbent manager's reputation, denoted by γ_t. At the initial date this is simply the reputation γ_0 of a manager picked from the pool. The investor then uses the reputation $\widehat{\gamma}_t$ from (9) and updates this prior after observing the manager's performance at $t + 1$. This yields a new reputation γ_{t+1}, which serves as the input into next period's maximization problem. The realization of γ_{t+1} is stochastic and depends on the manager's actual performance. When a newly employed manager acquires short-term information at date t, a performance observation at $t + 1$ consists of (s_t^{t+1}, d_{t+1}). Under long-term information acquisition it consists of $(s_t^{t+2}, p_{t+1}^{t+2})$. The transition function mapping $\widehat{\gamma}_t$ into γ_{t+1} thus depends on whether the manager acquired short-term or long-term information. Details of this function are provided in the proofs of Propositions 3 and 4, where the optimal policies are derived.

Moreover, the value function depends on a second state variable $I_t \in \{\varnothing, \{s_{t-1}^{t+1}\}\}$, which serves as a "memory" for the past signal. In some states of the world, the variable γ_t does not fully describe the state of the dynamic system. It is then necessary to introduce a further state variable that completes the description of the system's

---

[6] The multiplicity of equilibria in the trading game carries across to the overall equilibrium. However, this does not matter, because payoffs are uniquely determined across these equilibria, just as before.

[7] The equilibrium strategies are therefore Markov. Note, however, that this is not a restriction. It is the direct result of the simple structure of the problem.

history. This variable essentially serves as a record of the last signal when long-term information was acquired at the previous date. This is sometimes necessary, because the investor can usually only fully evaluate the manager *two* dates after he acquired long-term information. Suppose a manager acquires long-term information at $t - 1$. The investor can evaluate this manager's performance two periods later when the dividend d_{t+1} is realized. Knowing only the manager's reputation at date t is therefore generally not sufficient in order to update the belief about the manager at $t + 1$. In addition, the investor needs to remember the manager's date-$t - 1$ signal, which is recorded by $I_t = \{s_{t-1}^{t+1}\}$.

There are, however, some states when the manager has acquired long-term information at $t - 1$ that do allow a complete reputation update at date t. This happens if the investor learns at date t the dividend realization d_{t+1} from the price p_t^{t+1}. In that case she no longer needs to wait for the actual dividend realization in order to complete the belief update. This only happens when the manager receives $s_{t-1}^{t+1} = 0$, but the subsequent price reveals a nonzero dividend ($p_t^{t+1} \neq 0$). In this case we can set $I_t = \varnothing$. In all other cases of long-term information acquisition, either (i) the investor does not learn the dividend realization d_{t+1} at date t (when $p_t^{t+1} = 0$), or (ii) the price p_t^{t+1} reveals d_{t+1} due to the investor's own trade, which does not provide an independent confirmation of the future dividend payment. In these cases we set $I_t = \{s_{t-1}^{t+1}\}$.

Finally, when the manager acquires short-term information at date $t - 1$, the performance observation at date t is based on d_t and therefore contains all information necessary to perform a reputation update γ_t. We can then set $I_t = \varnothing$ whenever the manager acquired short-term information at $t - 1$. Formally, I_t is determined according to

$$(10) \qquad I_t = \begin{cases} \varnothing & \text{if } h_{t-1} = 1 \text{ or } (h_{t-1} = 2 \text{ and } s_{t-1}^{t+1} = 0, \, p_t^{t+1} \neq 0), \\ \{s_{t-1}^{t+1}\} & \text{otherwise.} \end{cases}$$

The investor then faces the following optimization problem:

$$(11) \qquad V(\gamma_t, I_t) = \max_{e_t, h_t} \left\{ E[\Pi(\widehat{\gamma}_t, h_t)] + \frac{1}{1+r} E[V(\gamma_{t+1}, I_{t+1})] \right\},$$

where the investor's instantaneous payoff $E[\Pi(\widehat{\gamma}_t, h_t)]$ can be calculated straightforwardly using (8) and (9), and $V(\gamma_{t+1}, I_{t+1})$ denotes the expected present value of all future payoffs under the optimal policy. Once the investor learns that a manager is a high type with probability one, learning stops and the optimal policy is to retain that manager and let him acquire long-term information at all future dates [this follows from the condition (6)].

REMARK 1 *If the condition* (6) *holds, the optimal policy at* $(\gamma_t = 1, I_t)$ *is* $e_T = 1$ *and* $h_T = 2$ *for all* $T \geq t$. *The value function is then given by*

$$(12) \qquad V(1) \equiv V(\gamma_t = 1, I_t) = \frac{(1 - \alpha) v_l}{4r(1+r)} \left(1 + \frac{1 - \beta}{2} \right).$$

5 Short-Term versus Long-Term Information Acquisition

This section presents the main result on the optimal choice of information horizon and employment policy. It proceeds in the following way. The state variables at the beginning of the game are $\gamma_0, \varnothing$; *i.e.*, the investor employs a manager from the pool, and the memory for past signals is empty. Starting from this point, the investor has the choice between short-term and long-term information acquisition. Firstly, I derive the optimal policy for all possible later realizations of the state variables starting from $\gamma_0, \varnothing$ under the assumption that the investor always chooses $h_t = 1$ when the state $(\gamma_0, \varnothing)$ occurs. I call this the short-term constrained optimal policy with associated value $V_s(\gamma_0, \varnothing)$ to the investor. The policy and its value are given in Proposition 3. I then provide the corresponding long-term constrained optimal policy and value $V_l(\gamma_0, \varnothing)$ under the alternative assumption that long-term information acquisition is chosen in state $(\gamma_0, \varnothing)$ (Proposition 4). It follows from the recursive structure of the problem that if it is optimal to induce short-term (long-term) information acquisition at any date when the state variable is $\gamma_0, \varnothing$, it will be optimal to do so whenever the state variable is $\gamma_0, \varnothing$. The value function is therefore simply given by $V(\gamma_0, \varnothing) = \max\{V_s(\gamma_0, \varnothing), V_l(\gamma_0, \varnothing)\}$. Note that the existence of an optimal solution is trivial here, because the problem is sufficiently simple to identify uniquely the optimal policy given that *either $h_t = 1$ or $h_t = 2$* is chosen in state $\gamma_0, \varnothing$.

5.1 Short-Term Information Acquisition

PROPOSITION 3 *Suppose the investor induces short-term information acquisition $(h_t = 1)$ in state $\gamma_t = \gamma_0, I_t = \varnothing$. The investor's short-term constrained optimal policy is then contingent on performance observations (s_t^{t+1}, d_{t+1}) as follows:*

(13)
$$\begin{aligned} e_{t+1} = 1, h_{t+1} = 2 & \quad \textit{if } s_t^{t+1} = d_{t+1} \in \{-1, 1\}, \\ e_{t+1} = 1, h_{t+1} = 1 & \quad \textit{if } s_t^{t+1} = d_{t+1} = 0, \\ e_{t+1} = 0, h_{t+1} = 1 & \quad \textit{otherwise.} \end{aligned}$$

The investor's expected discounted payoff from this policy is given by

(14)
$$V_s(\gamma_0, \varnothing) = \frac{(1 + r)\, E[\pi(\gamma_0, 1)] + \frac{1}{2}\gamma_0 \nu_s V(1)}{r + \frac{1}{2}\gamma_0 \nu_s}.$$

The short-term constrained optimal policy in this case takes a particularly simple form, which can be summarized by

(15)
$$\begin{aligned} e_t = 1, h_t = 2 & \quad \textit{if } \gamma_t = 1, \ I_t = \varnothing, \\ e_t = 1, h_t = 1 & \quad \textit{if } \gamma_t = \gamma_0, I_t = \varnothing, \\ e_t = 0, h_t = 1 & \quad \textit{if } \gamma_t < \gamma_0, I_t = \varnothing. \end{aligned}$$

If a newly employed manager performs badly once (he receives a zero signal when the dividend is nonzero), he is fired and replaced by a new manager from the pool. The new manager will be asked to acquire short-term information. If the manager

gets a zero signal and the dividend is zero, he is retained, but continues to trade on short-term information. If the manager trades in the correct direction, he is retained and asked to trade on long-term information. From Remark 1 we know that he will continue to acquire long-term information ever thereafter. This then also gives a complete description of an optimal policy under the assumption that a newly employed manager acquires short-term information.

Note that the investor is indifferent between retaining and firing the manager after the performance observation $s_t^{t+1} = d_{t+1} = 0$. Therefore, there exists another policy that yields the same payoff to the investor and in which a newly employed manager is always fired unless he generates a nonzero signal.

5.2 Long-Term Information Acquisition

When the manager acquires long-term information, we need to distinguish the cases (a) where he does so for the first time, and (b) where he does so repeatedly. In principle one may also consider the case where a manager goes back and forth between acquiring long-term and short-term information several times. However, as shown below, it is never optimal to induce short-term information acquisition directly after long-term information acquisition. Hence, whenever a manager acquires long-term information, it is either because he does so for the first time, or because he does so repeatedly without interruption. It follows that we need not consider the case of inducing short-term information unless $I_t = \varnothing$, which is the case covered in section 5.1.

LEMMA 2 *If it is optimal to let a newly employed manager acquire long-term information, it is never optimal to switch to short-term information acquisition at a later date.*

From Lemma 2 it follows that any policy for which long-term information acquisition is optimal in state $(\gamma_0, \varnothing)$ is fully characterized by considering the employment policy only. In the subsequent description the policy choice $h_t = 2$ is omitted for brevity.

PROPOSITION 4 *Suppose the investor induces long-term information acquisition $(h_t = 2)$ in state $\gamma_t = \gamma_0$, $I_t = \varnothing$. The following employment decisions are then long-term constrained optimal at date $t + 1$:*

$$
(16) \qquad \begin{aligned} e_{t+1} &= 1 \quad && \text{if } s_t^{t+2} \in \{-1, 1\}, \\ e_{t+1} &= 0 \quad && \text{otherwise.} \end{aligned}
$$

The investor's payoff in this case is given by

$$
(17) \qquad V_l(\gamma_0, \varnothing) = \frac{(1 + r)\, E[\pi(\gamma_0, 2)] + \frac{1}{2}\gamma_0 v_l V(1)}{r + \frac{1}{2}\gamma_0 v_l}.
$$

Again the employment policy is straightforward. If the manager receives a nonzero signal, he is retained, because only high types receive such a signal. From Remark 1

we know that this manager will now always be retained and asked to acquire long-term information. If the manager receives a zero signal, he is fired. The newly employed manager then falls under the rules of Proposition 4, *i.e.*, the state is $(\gamma_0, \varnothing)$, and he will acquire long-term information. A complete description of the long-term constrained optimal policy is therefore given by

(18)
$$e_t = 1, h_t = 2 \quad \text{if} \quad \gamma_t \geq \gamma_0, I_t \in \left\{\varnothing, \left\{s_{t-1}^{t+1}\right\}\right\},$$
$$e_t = 0, h_t = 2 \quad \text{if} \quad \gamma_t < \gamma_0, I_t \in \left\{\varnothing, \left\{s_{t-1}^{t+1}\right\}\right\}.$$

We are now in a position to identify the set of parameters for which the investor would prefer inducing short-term information acquisition with a newly employed manager. Since $V(\gamma_0, \varnothing) = \max\{V_s(\gamma_0, \varnothing), V_l(\gamma_0, \varnothing)\}$, short-term information acquisition is preferred when (14) is larger than (17). It is also straightforward to show that the intersection of the resulting parameter space with the restriction $v_s > v_l$ and (6) is nonempty. To see this most clearly, set $\beta = 1$, which implies that (6) is satisfied, *i.e.*, acquiring long-term information is first best. The comparison between $V_s(\gamma_0, \varnothing)$ and $V_l(\gamma_0, \varnothing)$ then yields the following condition:

COROLLARY 2 *Suppose $\beta = 1$. Then for*

(19)
$$\tfrac{1}{2}(1 - \alpha)(v_s - v_l) - \tfrac{1}{2}\gamma_0 v_s > r$$

the investor's equilibrium choice of information horizon for a newly employed manager is short-term.

If α is too large, the investor does not benefit sufficiently from identifying a high-type manager to be willing to incur the cost of doing so. The cost consists of the reduction in immediate trading profits that results from letting the manager acquire the less profitable short-term signal. Similarly, if r is large, the investor's opportunity cost of time becomes so high that she is no longer willing to forgo current profits in exchange for increased profits in the future. Finally, when γ_0 is small it is particularly costly to make the mistake of firing a high-type manager, because it is unlikely that another high type will be picked at the next date. This mistake occurs less frequently under short-term information acquisition.

Under either short-term or long-term information acquisition the investor learns that a manager is a high type only when the manager receives a nonzero signal. Since high types do not always receive a nonzero signal, they may go unrecognized (i) when the zero signal is accurate ($d_{t+h_t} = 0$), or (ii) when they happen not to learn that the future dividend is different from zero. The investor therefore makes the mistake of rejecting high types relatively often. Moreover, it is harder under long-term information acquisition to receive a signal that allows a manager to be identified as a high type because $v_l < v_s$. It follows that high types are rejected more frequently when long-term information is acquired. The investor may therefore prefer the manager to acquire short-term information.

An important difference between long-term and short-term information acquisition is the desirability of firing a manager. Under short-term information acquisition the investor strictly prefers to fire a manager only when he mistakenly receives the

zero signal, *i.e.*, when $s_t^{t+1} = 0$, $d_{t+1} \in \{-1, 1\}$. Under long-term information acquisition a manager is *always* fired when he does not trade. In particular, he is fired even when the subsequent price stays at zero, which provides some confirmation that the dividend payment is zero. The investor observes the price p_{t+1}^{t+2} as a *noisy* indicator of whether an arbitrage opportunity existed. When $p_{t+1}^{t+2} = 0$ the investor does not know whether this is an accurate reflection of the asset's true value or whether it simply indicates a failure of the market price to reflect the true value. Therefore the price $p_{t+1}^{t+2} = 0$ does not show conclusively to the investor whether her manager correctly received a zero signal or instead was unable to identify an arbitrage opportunity. Given this information, an observation of $s_t^{t+2} = 0$ is more likely to originate from a low type, and the manager's reputation therefore falls below γ_0. At this point it is strictly optimal to fire him. Hence, noisy prices carry a cost to the manager who acquires long-term information, because he faces a high likelihood of being fired. And they carry a cost to the investor, because she finds it harder to identify high-type managers as such.

6 *Incentive Compatibility*

So far it has been assumed that there is no agency problem between the investor and the manager. This could be the case, for example, because the signal generated by the manager is *hard* information, *i.e.*, the investor can verify it. In practice, some fund managers employ algorithms for generating trading recommendations. The hard information setup corresponds to the case where such a manager essentially hands the algorithm over to the investor, who can verify the resulting trading recommendations and choose the trades herself.

Suppose instead that the manager's signal is *soft*, *i.e.*, the investor cannot observe the manager's signal. In that case, the investor can learn from the manager's trade instead of the signal realization in much the same way, as long as he continues to trade according to the profit-maximizing strategies (1) and (2).

Note that the manager does not have an incentive to deviate from the equilibrium trading strategy if he receives a nonzero signal $s_t^{t+h_t} \neq 0$: His identity (high type) is immediately revealed to the investor by trading on the signal. This increases the manager's bargaining power and subsequent payoffs, so he clearly would not deviate. The manager, however, may wish to deviate by trading after $s_t^{t+h_t} = 0$. The potential benefit from doing so stems from (i) being thought of as a high type immediately and thereby increasing his bargaining power more quickly, and/or (ii) increasing his probability of being retained. Suppose also that the manager has limited liability (or no wealth), so that the agency problem cannot be resolved in a trivial way.[8]

When the manager acquires short-term information and receives a zero signal, his likelihood of being retained in the next period is highest if he does not trade, because $\text{prob}(d_{t+1} = 0|s_t^{t+1} = 0) > \text{prob}(d_{t+1} \neq 0|s_t^{t+1} = 0)$. The only potential benefit from

---

[8] With sufficient pledgeable wealth the investor could always impose a punishment for deviation that would resolve the agency problem at zero cost.

trading is therefore to be thought of as a high type immediately so as to increase his bargaining power more quickly. Compare this with long-term information acquisition, where the investor's employment policy is particularly harsh: a manager is fired unless he trades. The manager may therefore trade in order to increase his likelihood of being retained *and* in order to increase his bargaining power more quickly. This increases his overall incentive to deviate, and implementing the desired trading strategy becomes more costly under long-term information acquisition. The major effect of introducing an agency problem would therefore be to shift the balance in favour of short-term information acquisition.

This finding is consistent with VON THADDEN [1995], who shows that a long-term project may not be incentive-compatible when outside financiers make continuation of the firm contingent on short-term performance. In contrast to that, incentive compatibility under short-term information acquisition in my model is easier to satisfy, because the investor optimally applies a softer employment rule. This is because of the poorer information available under long-term information acquisition at the time of taking the decision to retain a manager.

Finally, the investor can reduce the cost of rendering the long-term trading strategy (1) incentive-compatible, by applying a softer employment rule: She could commit to retaining the manager if he does not trade and the subsequent price stays at zero. Interestingly, the cost of this policy depends on the information content of the stock price: If the quality of the short-term speculator's information, β, is high, then the stock price is more likely to reveal when a manager should have traded but failed to do so. When β is low (prices are not very informative), the investor is less likely to be able to detect such a failure, and this policy becomes too expensive. The investor therefore switches back to a price-insensitive firing policy. This clarifies the assumption made in SHLEIFER AND VISHNY [1997] regarding investors' fund-switching behaviour in response to interim price movements. We would expect investors to respond to interim performance observations only when prices may be informative. Noise traders may affect interim performance, but the more noisy prices become, the less sensitively one would expect investors to react to interim performance observations.

7   Conclusion

This paper addresses the question why fund managers may display a bias towards short-term information acquisition when it is more profitable to acquire and trade on long-term information instead. It shows that such short-termism may be driven by an investor's desire to learn about the unknown ability of a fund manager. The results are derived in a setting where the quality of the manager's information affects (i) the expected profits from trading on it, and (ii) the investor's ability to draw inferences about managerial ability from performance observations. When the quality of short-term information is higher than that of long-term information, the investor may prefer the former, since it allows her to draw sharper inferences about

the manager's ability. There is, however, a cost of learning, because short-term information is already partially disseminated into the price, which reduces profits from trading on it.

Appendix

A.1 Proof of Proposition 1

The investor's optimal trading strategy in asset τ can be determined by maximizing the expected discounted profit from trade in asset τ alone. This is because the investor's play in asset τ does not affect his own payoffs in other assets or the play by the market makers or the short-term speculators in other assets. Therefore, the trading strategy for each asset can be analyzed in isolation.

The following order flows may be observed by the market maker for asset $t+2$ under the trading strategy (1):

$Q_t^{t+2} = \{1, 1\}$: The investor's buy order is revealed; the market maker infers that the dividend is high ($d_{t+2} = 1$) and sets $p_t^{t+2} \left(Q_t^{t+2} = (1, 1) \right) = P^2$.

$Q_t^{t+2} \in \{\{1\}, \{-1\}\}$: The market maker knows that the investor did not submit an order, either because she received the signal $s_t^{t+2} = 0$, or because she did not acquire long-term information. Bayesian updating yields $p_t^{t+2} = 0$.

$Q_t^{t+2} = \{-1, 1\}$: The market maker knows that the investor submitted an order, but he cannot tell whether it was a buy or a sell order. Bayesian updating yields $\text{prob}(d_{t+2} = 1 | Q_t^{t+2} = \{-1, 1\}) = 1/2$. Similarly, $\text{prob}(d_{t+2} = -1 | Q_t^{t+2} = \{-1, 1\}) = 1/2$ and $\text{prob}(d_{t+2} = 0 | Q_t^{t+2} = \{-1, 1\}) = 0$. The price is therefore

$$p_t^{t+2} \left(Q_t^{t+2} = \{-1, 1\} \right) = \frac{1}{(1+r)^2} \left(1 \cdot \frac{1}{2} + (-1) \cdot \frac{1}{2} \right) = 0.$$

$Q_t^{t+2} = \{-1, -1\}$: Order flow reveals the investor's sell order, and the price fully reflects this information: $p_t^{t+2} \left(Q_t^{t+2} = \{-1, -1\} \right) = -P^2$.

The market maker's price-setting strategy, including off-the-path prices, is summarized below:

$$\text{(A1)} \quad p_t^{t+2} = \begin{cases} P^2 & \text{if } Q_t^{t+2} \in \{\{1, 1\}, \{x, 1\}, \{x, -1\}\}, \ x > 0, \ x \neq 1, \\ -P^2 & \text{if } Q_t^{t+2} \in \{\{-1, -1\}, \{x, 1\}, \{x, -1\}\}, \ x < 0, \ x \neq -1, \\ 0 & \text{otherwise.} \end{cases}$$

For off-the-equilibrium-path observations $Q_t^{t+2} \in \{\{x, 1\}, \{x, -1\}\}$ where $x \notin \{-1, 0, 1\}$, the market maker believes that a deviation $x > 0$ ($x < 0$) stems from a positively (negatively) informed investor and sets P^2 ($-P^2$).

Now consider optimality of the strategy (1). If the investor has a signal $s_t^{t+2} \in \{-1, 1\}$, she is clearly better off trading exactly -1, or 1, since in some states of the world her order is not revealed and she can make a profit. If the signal is $s_t^{t+2} = 0$, it is optimal not to trade at all, because if the investor were to trade (say buy), she would have to pay a high price $p_t^{t+2} > 0$ with positive probability (and never

$p_t^{t+2} < 0$). Given that her best estimate of the value of the stock conditional on her signal is zero, she would make a loss in expectation on such a trade.

Consider now order flow and prices for asset $t + 1$. Suppose that trade at date $t - 1$ did not reveal information, *i.e.*, $Q_{t-1}^{t+1} \in \{\{1\}, \{-1\}, \{-1, 1\}\}$. The price may move up or down when subsequent order flow does reveal private information. Orders reveal positive information when $Q_t^{t+1} \in Q^+ \equiv \{\{1, 1, 1\}, \{-1, 1, 1\}, \{1, 1\}\}$, in which case the price goes up to P. When $Q_t^{t+1} \in Q^- \equiv \{\{-1, -1-, 1\}, \{-1, -1, 1\}, \{-1, -1\}\}$, negative information is revealed and the price drops to $-P$.

Moreover, define the *set* of orders that can occur along the equilibrium path as $\widehat{Q}_{t-1}^{t+1}$ and $\widehat{Q}_t^{t+1} (Q_{t-1}^{t+1})$. The market maker's price-setting strategy can then be written as

$$
\text{(A2)} \quad p_t^{t+1} =
\begin{cases}
(1+r)p_{t-1}^{t+1} & \text{if } p_{t-1}^{t+1} \in \{-P^2, P^2\} \text{ and } Q_t^{t+1} \in \widehat{Q}_t^{t+1}, \\
0 & \text{if } p_{t-1}^{t+1} = 0 \text{ and } Q_t^{t+1} \in \{\{-1, 1\}, \{1\}, \{-1\}\}, \\
P & \text{if } Q_t^{t+1} \in Q^+, \text{ or } Q_t^{t+1} \notin \widehat{Q}_t^{t+1} \text{ and deviation } x > 0, \\
-P & \text{if } Q_t^{t+1} \in Q^-, \text{ or } Q_t^{t+1} \notin \widehat{Q}_t^{t+1} \text{ and deviation } x < 0.
\end{cases}
$$

The investor optimally trades again in asset $t + 1$ if her date-$t - 1$ signal was not revealed in p_{t-1}^{t+1}. Since there is a positive probability that her date-t trade will again be nonrevealing, she earns a strictly positive expected profit from doing so. The short-term speculator also earns a strictly positive expected profit from buying (selling) one unit after positive (negative) news. Again, let the market maker have the same off-the-equilibrium-path belief as before, so that a deviation $x > 0$ ($x < 0$) results in a price P ($-P$). This makes deviations from the equilibrium trading strategy unprofitable for the same reason as before. Also, by the same argument as before, the investor and the short-term speculator do not wish to trade if they have a zero signal, because it would move the price in an unfavourable direction with positive probability.

Consider next date t equilibrium prices and orders following $p_{t-1}^{t+1} \in \{-P^2, P^2\}$, *i.e.*, the date-$t - 1$ order revealed the investor's private information (or there was an off-the-path order). In the proposed equilibrium, orders following P^2 can be $\{1, 1, 1\}, \{-1, 1, 1\}, \{1, 1\}$, and $\{-1, 1\}$. The price is then set at $p_t^{t+1} = P$. Similarly, after $-P^2$, when orders can be $\{-1, -1, -1\}, \{1, -1, -1\}, \{-1, -1\}$, and $\{-1, 1\}$, the price is set at $p_t^{t+1} = -P$. Once again, one can apply the same off-the-path beliefs and prices: Denote by x the deviation from the above equilibrium order. The market maker then sets $p_t^{t+1} = P$ if $x > 0$ and $p_t^{t+1} = -P$ if $x < 0$. Since all information is already revealed in the price, the investor and the short-term speculator make exactly zero profits if they continue to trade. They are therefore indifferent between trading and not, and for the equilibrium presented here it is assumed that they continue to trade.

Note also that the equilibrium prices (on and off the path) do not allow for a profit-increasing deviation for more complex intertemporal strategies. Any attempt to drive down the price for asset $t + 2$ at t with a view to buying cheaply at $t + 1$ does not work: If the date-$t + 1$ buy order is $\neq 1$, then the price will immediately jump to P.

If it is 1, we need to distinguish three cases. Firstly, the date-t price was 0. In that case the probability of getting the zero price at $t + 1$ is unaffected by the initial trade, so it should be chosen independently. Secondly, the date-t price was negative because of an off-the-path order at t. In that case a buy order of 1 is off the path and will lead to a high price P. Finally, the date-t price was negative because of an on-the-path order at that date. This could only have happened because the investor submitted -1. This yields profits $-P - p_{t+1}^{t+2}$, which is negative according to Table 1. The same is true for the converse strategy of driving the price up at t in order to sell at a high price at $t + 1$.   Q.E.D.

A.2 Proof of Proposition 2

To calculate trading profits from long-term information acquisition, note that with probability $(1/2)(1 - \gamma_t v_l)$ the investor receives the signal $s_t^{t+2} = 0$ and therefore does not trade. With probability $(1/4)(\gamma_t v_l)$ the dividend payment is high *and* the investor learns it. She submits a buy order, which is not revealed in order flow (*i.e.*, $Q_t^{t+2} = \{-1, 1\}$ and $p_t^{t+2} = 0$) with probability 1/2. The same is true for the low dividend realization and the corresponding sell order. In all other states of the world the investor makes a zero profit. This yields expected (discounted) trading profits from the first-round trade of

$$\frac{1}{4}\gamma_t v_l \cdot \frac{1}{2}\left(\frac{1}{(1+r)^2} - 0\right) + \frac{1}{4}\gamma_t v_l \cdot \frac{1}{2} \cdot (-1)\left(-\frac{1}{(1+r)^2} - 0\right) = \frac{\frac{1}{4}\gamma_t v_l}{(1+r)^2}.$$

In the second period trade may be profitable if the first-round trade was not revealed in order flow. This happens with probability $(1/4)(\gamma_t v_l)$. The second-round trade only yields further profits if order flow Q_{t+1}^{t+2} is also not revealing. This happens when the short-term speculator receives the signal $\overline{s}_{t+1}^{t+2} = 0$ (probability $1 - \beta$) and the noise trader submits an order in the opposite direction to the investor's (probability 1/2). The expected discounted profits from the second trading round are therefore given by

$$\frac{\frac{1}{4}\gamma_t v_l}{(1+r)^2} \cdot \frac{1 - \beta}{2}.$$

Summing yields (4).

The investor makes a discounted trading profit of $1/(1 + r)$ when she acquires short-term information and receives signal $s_t^{t+1} \neq 0$ (probability $(1/2)(\gamma_t v_s)$) and the speculator does not submit an order (probability $1 - \beta$) and the noise trader submits an order in the opposite direction to the investor's (probability 1/2). Otherwise the investor's order is revealed to the market maker and the price adjusts to the true value of the asset, leaving no trading profit. This yields overall expected trading profit of

$$\frac{1}{1 + r} \cdot \frac{1}{2}\gamma_t v_s \cdot (1 - \beta) \cdot \frac{1}{2}.$$

Finally, in order to prove that profits are identical across all PBEs, we need to show uniqueness outside the trivially different equilibria. Consider first the possibility of equilibria in which an informed agent (investor or speculator) does not trade. At the

above efficient prices $p_t^{t+2}(\{-1, 1\}) = 0$ and $p_{t+1}^{t+2}(\{-1, 1\}) = 0$, an informed trader makes a strictly positive expected profit from trade. Moreover, in any equilibrium, price efficiency implies that a trader makes zero profits when his trade is revealed. No trade can therefore be an equilibrium only if the price fully reveals the trader's information, even when order flow is not fully revealing. This is impossible given the market maker's information set.

Consider now the possibility of trade in the opposite direction to the stipulated equilibrium: say an informed trader sells (buys) after receiving positive (negative) news. In that case price efficiency requires $p_t^{t+2}(\{-1, -1\}) = P^2$ (and similarly for date-1 trades). This, however, provides an incentive to deviate to the informed trader, because he would now make a strictly positive profit from selling after negative news if the price under nonrevealing order flow is 0. The only way to deter this would be to make the nonrevealing price negative. This, however, is only consistent with price efficiency if the trader is more likely to trade (sell) on positive news. This cannot be an equilibrium, because it would provide an incentive to buy after positive news. This argument also generalizes to the case where mixed strategies are played. The price would be $p_t^{t+2}(\{-1, -1\}) > -P^2$, which provides the manager with an incentive to deviate just as above.   *Q.E.D.*

A.3 Proof of Lemma 1

Firstly, note that the profit-maximizing trading strategy also maximizes the investor's expected payoff, because the latter is strictly increasing in profits. Moreover, the only long-run player is the investor. The market makers and short-term speculators only play during the one or two periods before the dividend date for their asset. They therefore choose their strategies so as to maximize expected profit (speculator) or implement price efficiency (market maker), *i.e.*, their objective functions are the same as in the trading game of Proposition 1.

The effect of the investor's trading strategy may be threefold. Firstly, it directly affects the trading profits. Secondly, it may affect learning about the manager's type, and finally it may affect future actions by later generations of short-term speculators and market makers. This last possibility has already been ruled out in the proof of Proposition 1. In order to prove the lemma, I show that the choice of trading strategy only has the first effect. From this it follows that the investor or manager chooses the trading strategy so as to maximize expected discounted trading profits given the signal realizations. All players therefore choose their strategies maximizing the same payoff functions as in Proposition 1. Given that their information sets are also unchanged, this implies that the set of equilibrium strategies must coincide.

Consider first how the investor's trading strategy affects learning about his type. The investor learns about the manager's type by comparing his signal realization with later realized dividends and prices. Under short-term information acquisition at date t, the investor updates her belief based on the observation of s_t^{t+1} and d_{t+1}. The trading strategy therefore has no effect on learning. Under long-term information acquisition at date t, a performance observation at $t + 1$ consists of $(s_t^{t+2}, p_{t+1}^{t+2})$.

When $s_t^{t+2} \in \{-1, 1\}$, the investor knows that the manager is a high type, and learning stops. Again, the trading strategy is irrelevant in that case. When $s_t^{t+2} = 0$, the future price may move up to P or down to $-P$, revealing that the manager could have received but failed to receive a signal. Since the price is only a noisy signal for the short-term speculator's information, the investor could learn more by inducing the short-term speculator to reveal his signal $\overline{s}_{t+1}^{t+2}$ directly. Since side payments are ruled out, the only way for the investor to achieve this would be to induce the speculator to trade so as to reveal his order to the market maker. This can be achieved if it is possible to compensate the speculator for the forgone trading profits by leaving him with higher trading profits in other states of the world. The only way for the investor to do that is by refraining from trading at date t when the manager has received a signal $s_t^{t+2} = -1$ or 1. In that case, however, the investor no longer needs to learn about the manager's ability, and therefore she would prefer to trade on such a signal. This strategy is therefore not time-consistent.   *Q.E.D.*

A.4 Proof of Proposition 3

The following performance observations at $t + 1$ following $h_t = 1$ in state $(\gamma_0, \varnothing)$ are possible.

(i) If $d_{t+1} \in \{-1, 1\}$, then with probability $\widehat{\gamma}_t \nu_s$ the manager observes a high signal $s_t^{t+1} = d_{t+1}$. Since only high types may get this signal, the posterior probability that the manager is a high type is equal to one:

$$\gamma_{t+1}(\widehat{\gamma}_t, d_{t+1} \in \{-1, 1\}, s_t^{t+1} = d_{t+1}) = 1.$$

The investor therefore switches to long-term information acquisition in the subsequent period and receives a present value of $V(1)$ (see Remark 1). This case occurs with probability $(1/2)(\widehat{\gamma}_t \nu_s)$.

(ii) If $d_{t+1} \in \{-1, 1\}$, the manager receives a signal $s_t^{t+1} = 0$ with probability $1 - \widehat{\gamma}_t \nu_s$. In this case the posterior can be calculated from Bayesian updating as

$$(A2) \qquad \gamma_{t+1}(\widehat{\gamma}_t, d_{t+1} \in \{-1, 1\}, s_t^{t+1} = 0) = \frac{\widehat{\gamma}_t (1 - \nu_s)}{1 - \widehat{\gamma}_t \nu_s},$$

which is smaller than $\widehat{\gamma}_t$. If $\widehat{\gamma}_t = \gamma_0$ the manager's reputation drops to below γ_0. It is then clearly optimal to fire the manager and hire a new one, who then acquires short-term information by assumption. The investor receives continuation payoff $V_s(\gamma_0, 0)$.

(iii) Finally, if $d_{t+1} = 0$ (probability 1/2), the high- and the low-type manager receive $s_t^{t+1} = 0$. This yields

$$\gamma_{t+1}(\widehat{\gamma}_t, d_{t+1} = 0, s_t^{t+1} = 0) = \widehat{\gamma}_t,$$

and the investor does not learn anything about the manager's type. When $\widehat{\gamma}_t = \gamma_0$ the investor receives $V_s(\gamma_0, \varnothing)$.

Hence, we can write

$$(A3) \quad V_s(\gamma_0, \varnothing) = E[\pi(\gamma_0, 1)] + \frac{1}{1+r}\left[\frac{1}{2}\gamma_0\nu_s V(1) + \left(1 - \frac{1}{2}\gamma_0\nu_s\right)V_s(\gamma_0, \varnothing)\right].$$

Solving (A3) yields (14). Note that there is a one-to-one relationship between the performance observations and the reputation updates. This allows us to formulate the optimal policy directly as a function of the performance observations, as is done in (13). *Q.E.D.*

A.5 Proof of Lemma 2

Suppose a newly employed manager acquired long-term information at $t-1$. If he receives a nonzero signal, the investor updates her belief about his ability to 1. By the assumption (6) it is then optimal for him to acquire long-term information forever in the future. If the manager receives the signal $s_{t-1}^{t+1} = 0$, his reputation falls to $\gamma_t (s_{t-1}^{t+1} = 0) = [\gamma_0 (1 - v_l)]/[1 - \gamma_0 v_l]$. Moreover, the investor updates her belief that $d_{t+1} \neq 0$ to $(1 - \gamma_0 v_l)/(2 - \gamma_0 v_l) < 1/2$. If the investor asks the manager to acquire short-term information at t, this yields an expected payoff of

$$V_s \left(\gamma_t \left(q_{t-1}^{t+1} = 0 \right), \{0\} \right) = \frac{1 - \gamma_0 v_l}{2 - \gamma_0 v_l} \frac{\frac{\gamma_0 (1 - v_l)}{1 - \gamma_0 v_l} v_s \left(\frac{1-\beta}{2} + V(1) \right)}{1 + r}$$

$$+ \frac{1 - \gamma_0 v_l}{2 - \gamma_0 v_l} \frac{\left(1 - \frac{\gamma_0 (1 - v_l)}{1 - \gamma_0 v_l} v_s \right)}{1 + r} V(\gamma_0, \varnothing)$$

$$+ \frac{1}{2 - \gamma_0 v_l} \frac{V(\gamma_0, \varnothing)}{1 + r}.$$

Instead of retaining the manager and letting him acquire short-term information, the investor could have hired a new manager and let him acquire long-term information. The payoff from doing so would be $V_l (\gamma_0, \varnothing)$. If short-term information acquisition is optimal at t, we must have $V_s \left(\gamma_t \left(s_{t-1}^{t+1} = 0 \right), \{0\} \right) > V_l (\gamma_0, \varnothing)$. Moreover, if it was optimal to choose long-term information acquisition initially, we know that $V_l (\gamma_0, \varnothing) > V_s (\gamma_0, \varnothing)$. Hence, we would have $V_s \left(\gamma_t \left(s_{t-1}^{t+1} = 0 \right), \{0\} \right) > V_s (\gamma_0, \varnothing)$. Using (A3), it is straightforward to show that this leads to a contradiction. *Q.E.D.*

A.6 Proof of Proposition 4

Suppose a manager of reputation $\widehat{\gamma}_t$ acquires the long-term signal for the first time ($I_t = \varnothing$). Three performance observations are possible.

(i) The dividend payment is $d_{t+2} \in \{-1, 1\}$, and the manager receives $s_t^{t+2} = d_{t+2}$. The investor then updates her belief about the manager to $\gamma_{t+1} = 1$, retains the manager, and lets him acquire long-term information in the subsequent period (see Remark 1).

(ii) The manager receives a zero signal $s_t^{t+2} = 0$, but the subsequent price indicates that the dividend payment is nonzero: $s_t^{t+2} = 0$, $p_{t+1}^{t+2} \in \{-P, P\}$. This occurs when the dividend payment is nonzero (probability $1/2$), the manager does not learn this (probability $1 - \widehat{\gamma}_t v_l$), the short-term speculator receives signal $\overline{s}_{t+1}^{t+2} = d_{t+2}$

(probability β), and his order is revealed to the market maker (probability $1/2$). The investor therefore knows that $d_{t+2} \neq 0$, and his Bayesian update about managerial ability will not change at $t + 2$. The reputation at $t + 1$ changes to $\gamma_{t+1} = [\widehat{\gamma}_t (1 - v_l)]/[1 - \widehat{\gamma}_t v_l]$, and it is a sufficient statistic for the manager's performance history. We can therefore set $I_{t+1} = \varnothing$. Moreover, since $[\widehat{\gamma}_t (1 - v_l)]/[1 - \widehat{\gamma}_t v_l] < \widehat{\gamma}_t$ and the investor does not benefit from waiting to observe d_{t+2}, it is clearly optimal to fire a manager who started with reputation $\widehat{\gamma}_t = \gamma_0$ after this performance observation.

(iii) The manager receives a zero signal and the subsequent price stays at zero $\left(q_t^{t+2} = 0, p_{t+1}^{t+2} = 0\right)$. This can happen because the future dividend payment is zero (probability $1/2$), or because the dividend payment is nonzero and the manager received a signal $s_t^{t+2} = 0$ [probability $(1/2)(1 - \widehat{\gamma}_t v_l)$], and order flow is not revealing, either because the short-term speculator also did not trade (probability $1 - \beta$), or because his order was not revealed (probability $\beta/2$). The manager's reputation then drops to

$$\gamma_{t+1}^- = \frac{\widehat{\gamma}_t \left[1 + (1 - v_l)\left(1 - \frac{\beta}{2}\right)\right]}{1 + \eta_t},$$

where $\eta_t \equiv [1 - \widehat{\gamma}_t v_l][1 - (\beta/2)]$.

Note that a manager who started with reputation $\widehat{\gamma}_t = \gamma_0$ now has a reputation $\gamma_{t+1}^- < \gamma_0$. An investor can improve the expected ability of the manager employed by hiring a new one. On the other hand, when the manager is retained, the memory for the past trade is activated ($I_{t+1} = \{0\}$). From Lemma 2 we know that he will acquire long-term information again in the next period. Note that the investor learns about the manager from the date-$t + 1$ information acquisition in the same way, independently of whether the incumbent or a new manager acquires information. Hence, the rate of learning differs only due to the date-t information acquisition that a new manager does not have. The highest reputation that observing $(s_t^{t+2} = 0, d_{t+2})$ can generate occurs when the investor learns that the manager's zero signal at date t was in fact optimal, because $d_{t+2} = 0$. This yields an updated reputation $\gamma_{t+2}(\widehat{\gamma}_t, d_{t+2} = 0, s_t^{t+2} = 0) = \widehat{\gamma}_t$. Since we started with $\widehat{\gamma}_t = \gamma_0$, the investor can therefore do no worse by picking a manager from the pool immediately instead of waiting for a period. It follows that it is optimal to fire the incumbent manager after $\left(s_t^{t+2} = 0, p_{t+1}^{t+2} = 0\right)$.

Hence, a newly employed manager is only retained when he receives a nonzero signal and is thus identified as a high type. This happens with probability $(1/2)(\gamma_0 v_l)$ and yields future expected payoffs of $V(1)$. In all other states the manager is fired, yielding payoff $V_l(\gamma_0, \varnothing)$.

We can then write

$$(A4) \quad V_l(\gamma_0, \varnothing) = E[\pi(\gamma_0, 2)] + \frac{1}{1 + r}\left[\frac{1}{2}\gamma_0 v_l V(1) + \left(1 - \frac{1}{2}\gamma_0 v_l\right)V(\gamma_0, \varnothing)\right].$$

Solving (A4) yields (17). Q.E.D.

References

ALLEN, F., AND G. GORTON [1993], "Churning Bubbles," *Review of Economic Studies*, 60, 813–836.

– –, S. MORRIS, AND H. S. SHIN [2004], "Beauty Contests and Iterated Expectations in Asset Markets," unpublished Mimeo, London School of Economics.

BHATTACHARYA, S., AND P. PFLEIDERER [1985], "Delegated Portfolio Management," *Journal of Economic Theory*, 36, 1–25.

CHEVALIER, J., AND G. D. ELLISON [1997], "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 105, 1167–1200.

– – AND – – [1999], "Career Concerns of Mutual Fund Managers," *Quarterly Journal of Economics*, 389–432.

DASGUPTA, A., AND A. PRAT [2003], "Trading Volume with Career Concerns," unpublished Working Paper, London School of Economics.

DEMIRAG, I. [1995], "Short-Term Performance Pressures: Is There a Consensus View?" *European Journal of Finance*, 1, 41–56.

DOW, J., AND G. GORTON [1994], "Arbitrage Chains," *Journal of Finance*, 49, 819–849.

– – AND – – [1997], "Noise Trading, Delegated Portfolio Management, and Economic Welfare," *Journal of Political Economy*, 105, 1024–1050.

GÜMBEL, A. [2005], "Should Short-Term Speculators be Taxed, or Subsidised?" *Annals of Finance*, 1, 327–348.

HEINKEL, R., AND N. M. STOUGHTON [1994], "The Dynamics of Portfolio Management Contracts," *Review of Financial Studies*, 7, 351–387.

HOLDEN, C. W., AND A. SUBRAHMANYAM [1996], "Risk Aversion, Liquidity, and Endogenous Short Horizons," *Review of Financial Studies*, 9, 691–722.

KYLE, A. [1985], "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

MARSH, P. [1990], *Short-Termism on Trial*, Institutional Fund Managers' Association: London.

SHLEIFER, A., AND R. W. VISHNY [1990], "Equilibrium Short Horizons of Investors and Firms," *American Economic Review: Papers and Proceedings*, 80, 148–153.

– – AND – – [1997], "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.

SPIEGEL, M., AND A. SUBRAHMANYAM [1992], "Informed Speculation and Hedging in a Non-Competitive Securities Market," *Review of Financial Studies*, 5, 307–329.

STEIN, J. C. [1989], "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics*, 104, 655–669.

STOUGHTON, N. M. [1993], "Moral Hazard and the Portfolio Management Problem," *Journal of Finance*, 47, 2009–2028.

VON THADDEN, E. [1995], "Long-Term Contracts, Short-Term Investment and Monitoring," *Review of Economic Studies*, 62, 557–575.

Alexander Gümbel
Saïd Business School and Lincoln College
University of Oxford
Oxford, OX1 3DR
United Kingdom
E-mail:
alexander.guembel@sbs.ox.ac.uk