



Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning

Author(s): Canice Prendergast and Lars Stole

Source: *The Journal of Political Economy*, Vol. 104, No. 6 (Dec., 1996), pp. 1105-1134

Published by: The University of Chicago Press

Stable URL: <http://www.jstor.org/stable/2138935>

Accessed: 13/07/2009 17:47

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=ucpress>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Political Economy*.

Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning

Canice Prendergast

University of Chicago and National Bureau of Economic Research

Lars Stole

University of Chicago

This paper examines individual decision making when decisions reflect on people's ability to learn. We address this problem in the context of a manager making investment decisions on a project over time. We show that in an effort to appear as a fast learner, the manager will exaggerate his own information; but ultimately, he becomes too conservative, being unwilling to change his investments on the basis of new information. Our results arise purely from learning about competence rather than concavity or convexity of the rewards functions. We relate our results to the existing psychology literature concerning cognitive dissonance reduction.

I. Introduction

The decisions that we make reflect our competence. In the absence of immediate reliable information on their impact, the decisions

We are grateful to In-Koo Cho, Chip Heath, Owen Lamont, Eric Rasmusen, Sherwin Rosen, an anonymous referee, and especially Kent Daniel for helpful discussions and to Vijay Fafat for research assistance. In addition, this paper has benefited from the input of seminar participants at the California Institute of Technology, Dartmouth College, the London School of Economics, Michigan State University, the National Bureau of Economic Research, Northwestern University, Princeton University, Tel Aviv University, Texas A&M, the University of Chicago, University of Florida, University of Maryland, University of Michigan, University of Rochester, University of Wisconsin, and the 1995 World Congress of the Econometric Society. We are thankful for financial support from the Center for the Study of the Economy and the State, the Centel Foundation/Robert P. Reuss Faculty Research Fund, the William Ladany Research Fund, a Sloan Foundation Fellowship, and a National Science Foundation Presidential Faculty Fellowship. Any errors are our own.

[*Journal of Political Economy*, 1996, vol. 104, no. 6]

© 1996 by The University of Chicago. All rights reserved. 0022-3808/96/0406-0007\$01.50

themselves may tell much about the quality of a decision maker. For instance, a politician who frequently changes his position on some matter may be perceived as unsure of himself, whereas a new executive who exactly mimics the behavior of her predecessor may be felt to have no ideas of her own. This paper addresses how such inferences affect the decisions that individuals make.

Our primary concern is to understand how individuals change their behavior on receipt of new information. In an ideal world, every individual would behave like a rational Bayesian, optimally learning about the economic environment by correctly combining new information with prior knowledge and then using this information to maximize value. We depart from this ideal by assuming that individuals wish to acquire a reputation for quickly learning the correct course of action. Two possible distortions are identified: (i) exaggeration, where individuals respond too much to new information, and (ii) conservatism, where individuals do not change their behavior enough in light of new events. The principal result of the paper is that initially individuals will overreact to new information but after some period of time become unwilling to respond to new information suggesting that their previous behavior was wrong. We provide simple intuitive conditions outlining when each type of distortion is likely to occur. These phenomena are addressed in the context of a manager choosing investments on a project over time on the basis of private observations of the project's profitability.

A couple of different interpretations can be placed on our results. First, the reluctance of individuals with long tenure on a project to change their previous positions generates the "jaded old-timers" of the paper's title. However, while replacing these managers with new blood may induce more action, these new managers are likely to distort in the opposite direction, by exaggerating their own opinion. In this sense, they are the "impetuous youngsters," keen to show that they have ideas of their own.

An alternative interpretation concerns decision-making heuristics studied by social psychologists concerned with biases related to cognitive dissonance reduction. Two prominent biases that are often highlighted are (i) the sunk cost fallacy and (ii) the base rate fallacy. The sunk cost fallacy describes an unwillingness of individuals to respond to new information and corresponds to the notion of conservatism used here. The base rate fallacy or overconfidence effect concerns situations in which individuals overweight their own information to the detriment of information on the population and corresponds closely to the concept of exaggeration described above. Our results can then easily be recast in terms of these behavioral rules in which both phenomena derive from the *same* process of learning.

These results emerge from a model in which a manager chooses investments on a project in each of T periods. The project has unknown profitability, but the manager receives a private signal of its return in each period, representing the manager's own subjective estimate of the project's profitability. Managers differ in their ability: talented managers receive more precise estimates of underlying profitability than their less able counterparts and are therefore able to learn about and respond to the economic environment more quickly. This ability is assumed to be private information to the manager, but rents accrue to a manager if he is perceived to be talented. In particular, we assume that the preferences of the manager are linearly increasing in current profits and the end-of-period market perception of ability.

At the start of each period, the manager receives a new private observation on the project's return. With this information he constructs an estimate of profitability and chooses an investment appropriately. In response to this investment, the market updates its opinion on the manager's ability; hence decisions themselves reveal competence. Throughout the paper, inferences are drawn from the difference between actual investment and expected investment, so that the *variance* of the posterior estimate on profitability plays a key role in the market's inference of the manager's ability to learn.

To understand our results, consider the first period of the manager's tenure, when he draws his first observation on profitability. When choosing investments, talented managers trust their own opinion more than their less able counterparts. This implies that, from an *ex ante* perspective, efficient investment is more variable for talented managers than for untalented managers; thus the variance of the posterior is increasing in managerial ability. But then a manager will *exaggerate* his true beliefs about profitability on the margin to appear talented. If his new information suggests that marginal investment is worse than the prior, he scales back operations more than is efficient; if the information regarding investment is good, he places too much capital in these projects.¹ Hence the initial response of managers is to "go out on a limb," and the variance of investment exceeds its first-best level.

The key distinction between the initial period and future periods is that the manager has already made previous investments. There-

¹ This is consistent with empirical evidence cited in Staw and Ross (1989), where managers were rated highest by subjects in experiments in which those managers had staked out an extreme position and were persistent, yet were successful. Staw and Ross interpret this result as a "hero effect," giving special praise for managers who "stick to their guns" in the face of opposition.

fore, any decisions that he makes subsequently are an indication of the quality of both his previous and current information. This gives rise to two conflicting effects. First, the manager may wish to make large revisions to investment to show how much he trusts his most recent observation, as in period 1. Second, if he were really talented, he would already have been close to choosing efficient investment so that a truly talented manager should change little *since he has no need to*. This effect leads to too little variability in investment.

Whether the manager exaggerates his position or invests too conservatively depends on the relative size of the two effects. A feature of the paper is that it provides a clear interpretation of whether exaggeration and conservatism will arise beyond the first period. If the task carried out is sufficiently routine (in the sense that the initial contribution of the manager is great), the manager "should know what is going on" after only a few observations. In this case, the manager shows excessive reliance on his previous position and does not admit mistakes, so that the variance of investment is below its first-best level. If, however, learning occurs slowly, the manager continues to exaggerate his information. In addition, we also demonstrate that managers must act conservatively beyond some specific date.

The reason that the manager ultimately becomes conservative is that the investment decisions increasingly become associated with his previous contributions over time; initially, he has made little input to the project, but eventually changing the investment becomes associated with his previous errors. The extent of this result is addressed in Section VI, where we consider situations in which this feature of increased responsibility for errors may not be true. In particular, we consider different economic environments to provide additional comparative statics on our insights. First, we allow the true profitability of the project to change over time, where profitability follows a random walk. Here we show that conservatism becomes less likely, and indeed if the environment changes rapidly enough, the manager exaggerates forever. Therefore, the paper establishes a connection between a stable economic environment and the prevalence of conservatism. Second, we consider a case in which other information on profitability is available but the quality of that information does not depend on the manager. For example, lagged revenues could be observed. Here we show two results. First, the manager always acts conservatively with respect to this public information. Therefore, exaggeration of information can occur only if the quality of the information itself reflects the manager's talent. Second, the manager can ultimately become conservative in response to his own private information only if the quality of the best manager's information is better

than the quality of the other public signal. Otherwise, the manager exaggerates forever.

Section II describes the model and a few statistical results regarding normal learning over variances. In Section III, we illustrate the incentive of the manager to initially exaggerate, and Section IV considers the complementary case of conservatism. In Section V, we extend the model to examine the dynamic implications of reputation on exaggeration and conservatism. Section VI considers additional comparative statics and the robustness of our conclusions by examining extensions. Section VII reinterprets our results in the context of the behavioral decision-making literature and the recent literature in economics on reputation and herding by Scharfstein and Stein (1990) and Zwiebel (1995).

II. The Model

Consider a manager who commences a project in period 1 that ends in period $T \geq 1$. New investment, I_t , is publicly chosen each period by the manager after privately observing additional information. The profit per period from investment is linear in an unknown parameter μ and concave in investment; investments fully depreciate at the end of each period. Negative investment, $I_t < 0$, is interpreted as investment in a project negatively correlated with the return μ . Specifically, we suppose that profitability is

$$\pi_t = \mu I_t - \frac{1}{2} I_t^2,$$

where μ is a productivity parameter that is symmetrically unknown before period 1, but whose distribution is commonly known to be normal with mean zero and variance τ^2 . We assume throughout that the return to investment is specific to the firm; that is, the firm has complete property rights in the investment project. The full-information level of profits is obtained by setting $I_t = \mu$.

The manager receives an imperfect private observation of μ at the beginning of each period given by $m_t = \mu + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$. The talent of the manager lies in identifying the quality of projects and varies across managers. Managerial ability is characterized by σ , with low- σ managers being deemed high-ability. Ability σ is specific to the manager and not to the firm. As a consequence, a manager with a reputation for high ability will obtain rents. We assume that ability is private information to the manager but is commonly known at date 0 to be distributed according to a distribution F with density $f > 0$ and support $[\underline{\sigma}, \bar{\sigma}] \subset (0, \infty)$. We let $\hat{\mu}_t$ denote the manager's optimal posterior estimate of μ , conditional on his signals and under-

lying variance:

$$\hat{\mu}_t \equiv E[\mu | m_t, \hat{\mu}_{t-1}, \sigma],$$

where we need to condition only on $\hat{\mu}_{t-1}$, which is sufficient for $\{m_{t-1}, \dots, m_1\}$.

If the manager was rewarded solely on the basis of profit, he would have perfect incentives and choose $I_t = \hat{\mu}_t$ in each period. We assume that, in addition, the manager also cares about his reputation for identifying profitability. Ideally, we would allow the manager to make investments taking account of how that investment affects not only perceptions today but also all future possibilities. However, it is well known that multiperiod incentive and learning models are inherently nonstationary, where the incentive to take an action depends on the history of previous actions and (typically more problematic) the set of actions that might be taken in the future. This often makes it impossible for such models to generate clear insights without some simplifying assumptions.² Accordingly, we also place restrictions on preferences to simplify our analysis. Let Ω^{t-1} be all public information available at the start of period t . We assume that the manager's objective in period t is to maximize

$$V_t = \hat{\mu}_t I_t - 1/2 I_t^2 - \lambda E^*[\sigma | \Omega_{t-1}, I_t],$$

where $\lambda > 0$, and $E^*[\sigma | \Omega_{t-1}, I_t]$ is the market's *equilibrium* expectation of the manager's "inability" after he observes the historical information Ω_{t-1} and the current investment I_t .³ Thus our characterization of managerial preferences assumes that the manager cares only about *current* profits and his immediate end-of-period reputation.⁴ It is important that the manager cares about profits and not merely reputation; without preferences over profits, no separation can occur.

We are primarily interested in how the decisions that managers make (investments) affect how they are perceived. Consequently, for now we simplify the analysis and assume that all the updating occurs through the investments that the manager makes rather than through

² For example, Holmström and Milgrom (1987) assume the absence of income effects to generate the linearity of optimal contracts. Similarly, typical learning models such as those by Harris and Holmström (1982), Holmström (1982), and Gibbons and Murphy (1992) assume normality of errors plus uncertainty only over the mean of the parameters of interest, implying incentives that do not depend on previous actions.

³ Ability is assumed to matter since the manager could solicit a job elsewhere during the period, with his wage offer depending on perceived ability, as in Harris and Holmström (1982) and Holmström (1982). Our reduced form captures this phenomenon.

⁴ It is purely for notational simplicity that we assume that the manager's preferences are linear and time-independent in expected σ . Our results continue to hold more generally if the manager's welfare is strictly decreasing in expected σ for every t , provided that the return to perceived σ is independent of current profits.

any measures of profits. As a modeling exercise, we assume that π_t is essentially uninformative, so as to avoid the inference problems inherent in incorporating additional signals.⁵ The case of observable profits is considered in Section VI.

A. *The Relationship between Ability and the Distribution of Posteriors*

To understand the manager’s incentives, it is useful to begin by deriving the distribution of $\hat{\mu}_t$ for a manager with a given σ . The manager’s estimate of the profitability of the project during period 1 (after m_1 is observed) is given by

$$\hat{\mu}_1 = E[\mu|m_1, \sigma] = \frac{\tau^2}{\tau^2 + \sigma^2} m_1.$$

From DeGroot (1970, chap. 9), the posterior distribution of μ is normal with mean $\hat{\mu}_1$ and variance $V[\mu|\hat{\mu}_1, \sigma] = \tau^2 \sigma^2 / (\tau^2 + \sigma^2) < \tau^2$.

The ex ante variance of the type σ manager’s period 1 posterior is given by $V[\hat{\mu}_1|\mu_0, \sigma] = \tau^4 / (\tau^2 + \sigma^2)$. A lower σ implies a higher variation of the posterior about the market’s prior, $\mu = 0$. This arises because managers who have very precise information (low σ^2) will downplay the prior (high τ^2) and place more weight on their own information. Since the labor market does not observe m_1 , this will be the distribution that it will use to infer the ability of the manager.

Consider now later periods. We derive $V[\hat{\mu}_t|\hat{\mu}_{t-1}, \sigma]$ to determine whether the positive relationship between variance and ability described above holds beyond period 1. For all periods $t > 1$, optimal updating yields the result that μ is distributed normally with mean

$$\hat{\mu}_t = \left(\frac{\tau^2}{t\tau^2 + \sigma^2} \right) m_t + \left[\frac{(t-1)\tau^2 + \sigma^2}{t\tau^2 + \sigma^2} \right] \hat{\mu}_{t-1}$$

and variance $V[\mu|\hat{\mu}_t, \sigma] = \sigma^2 \tau^2 / (t\tau^2 + \sigma^2)$. Taking the variance of the posterior in the equation above, conditioning on $\hat{\mu}_{t-1}$ and σ , we have

$$V[\hat{\mu}_t|\hat{\mu}_{t-1}, \sigma] = \left(\frac{\tau^2}{t\tau^2 + \sigma^2} \right)^2 V[m_t|\hat{\mu}_{t-1}, \sigma]. \tag{1}$$

⁵ Formally, these results can be derived from a case in which a noisy inference is made on profits but the precision of the estimate on profitability goes to zero. This renders the model much simpler to solve without altering the flavor of the results. An alternative interpretation of our results is that there is random monitoring, and μ is determined if monitoring occurs and the resulting compensation following an audit is linear in profits. The outcome of the game is then the one that occurs in all states in which monitoring does not arise.

Equation (1) is the key to understanding the results that follow. The variance of the posterior, $V[\hat{\mu}_t|\hat{\mu}_{t-1}, \sigma]$, has two components that act in conflicting ways. First, the coefficient $\tau^2/(t\tau^2 + \sigma^2)$ measures the “weight” placed on the most recent observation. This is *decreasing* in σ for the reason that a lower value of σ implies that the current observation is trustworthy and so should considerably change the posterior. Therefore, all other things equal, the manager would like to give the impression that this coefficient is large to indicate that the current observation is valuable. This effect leads to excess variability in investment. In contrast, $V[m_t|\hat{\mu}_{t-1}, \sigma]$ identifies the “variability” of the current observation. This is *increasing* in σ because as σ rises (i) the previous observation did not yield a precise measure of true profitability and (ii) the noise associated with the current observation is large. The manager would like to give the impression that $V[m_t|\hat{\mu}_{t-1}, \sigma]$ is small, so that he had already “gotten it right.” This argues against excess variability in investment. As will be seen below, these two countervailing effects generate our findings.

Because $m_t = \mu + \epsilon_t$, we know that

$$V[m_t|\hat{\mu}_{t-1}, \sigma] = V[\mu|\hat{\mu}_{t-1}, \sigma] + \sigma^2 = \frac{\sigma^2\tau^2}{(t-1)\tau^2 + \sigma^2} + \sigma^2.$$

Substituting and simplifying, we have the conditional variance of $\hat{\mu}_t$.

LEMMA 1.

$$\bar{\sigma}_{\hat{\mu}_t}^2(\sigma) \equiv V[\hat{\mu}_t|\hat{\mu}_{t-1}, \sigma] = \frac{\sigma^2\tau^4}{(t\tau^2 + \sigma^2)[(t-1)\tau^2 + \sigma^2]},$$

where $\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)$ is understood to depend on $\hat{\mu}_{t-1}$.

Our results will largely be cast in terms of $\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)$, which identifies the variance of the posterior around the prior, $\hat{\mu}_{t-1}$. The key condition determining whether exaggeration or conservatism arises is whether this variance is increasing or decreasing in σ . In other words, does varying investment from the prior illustrate talent or a lack of talent? In the one-period case, we saw that $d\bar{\sigma}_{\hat{\mu}_1}^2(\sigma)/d\sigma < 0$, so that the posteriors of high-precision managers should be more variable than those of their low-precision counterparts. More generally, this relationship does not hold as

$$\frac{d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)}{d\sigma} < 0 \quad \text{iff} \quad \frac{\sigma^2}{\tau^2} > \sqrt{t(t-1)}. \tag{2}$$

Therefore, how opinions on ability are revised depends on the time period: If and only if $\sigma^2/\tau^2 > \sqrt{t(t-1)}$ does the variance of $\hat{\mu}_t$ increase with ability. Note that for $t = 1$ this condition holds for any σ , but for large enough t it is always violated. Therefore, the nature of

the distortion will depend on the manager’s tenure on the investment project.

Because the variance of the posterior depends on σ , inferences on ability can be drawn from the choice of investment. In fact, a specific relationship holds as demonstrated in the following lemma. All proofs are contained in the Appendix.

LEMMA 2. Suppose that $\hat{\mu}_t \sim \mathcal{N}(\hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma))$. If $\hat{\mu}_t = \hat{\mu}_{t-1}$, then $\partial E[\sigma | \hat{\mu}_t, \dots, \hat{\mu}_0] / \partial \hat{\mu}_t = 0$. If $\hat{\mu}_t \neq \hat{\mu}_{t-1}$ and $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma) / d\sigma > (<) 0$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$, then

$$\left(\frac{\partial E[\sigma | \hat{\mu}_t, \dots, \hat{\mu}_0]}{\partial \hat{\mu}_t} \right) (\hat{\mu}_t - \hat{\mu}_{t-1}) > (<) 0.$$

This lemma simply shows that if the variance of the posterior is increasing (decreasing) in ability, then larger (smaller) changes in investment reflect well on the manager’s ability.

B. Equilibrium Separation

The market observes investment in each period but does not directly see the posterior of the manager. Nonetheless, the manager may signal his posterior $\hat{\mu}_t$ to the market in order to indicate high ability. We show that indeed this conjecture is correct under some mild conditions. Along these lines, we concentrate our attention on separating equilibria.

First, we consider what can be signaled to the market via investment. It is immediate that in any equilibrium, separating or otherwise, the manager’s choice of investment will directly depend only on historical observables, Ω^{t-1} , and the actual posterior, $\hat{\mu}_t$; σ is relevant only insofar as it determines $\hat{\mu}_t$. The immediate result of this is that *the most informative signaling equilibria can at most separate managers of different posteriors, $\hat{\mu}_t$, for any given history*. Because we concentrate on fully separating equilibria, the market’s sequence of posterior inferences is a sufficient statistic for Ω^{t-1} . Thus we can denote histories by $h^{t-1} \equiv \{\hat{\mu}_s^*\}_{s=0}^{t-1}$ (where $\hat{\mu}_s^*$ is the believed posterior of the manager in period s) and focus on the expectation of σ conditional on the market’s inference h^{t-1} (which is correct in equilibrium). In a separating equilibrium, the market will perfectly infer the manager’s posterior from the current investment level, I_t , and its previous inferences, h^{t-1} . We let $\hat{\mu}_t^*(I_t, h^{t-1})$ represent this equilibrium market inference of the manager’s posterior so that the relevant market expectation of σ following an inferred history of h^{t-1} is $E[\sigma | \hat{\mu}_t^*(I_t, h^{t-1}), h^{t-1}]$. As such, the manager (at time t) maximizes

$$\max_I V_t(I_t, \hat{\mu}_t, h^{t-1}) \equiv \hat{\mu}_t I_t - \frac{1}{2} I_t^2 - \lambda E[\sigma | \hat{\mu}_t^*(I_t, h^{t-1}), h^{t-1}]. \quad (3)$$

Solving this program, one obtains a set of conditions that can be shown to be necessary and sufficient for a fully separating perfect Bayesian-Nash equilibrium (PBE). It is easiest to first state our definition of a separating PBE and then characterize a simpler set of equivalent conditions.

DEFINITION. The functions $\{I_t^*(\hat{\mu}_t, h^{t-1}), \hat{\mu}_t^*(I_t, h^{t-1})\}_{t=1}^T$ represent a fully separating PBE iff, for all t and h^{t-1} , (1) $I_t^*(\hat{\mu}_t, h^{t-1}) \in \text{argmax}_{I \in \mathbb{R}} V_t(I, \hat{\mu}_t, h^{t-1})$ for all $\hat{\mu}_t \in \mathbb{R}$ and (2) $\hat{\mu}_t^*(I_t, h^{t-1})$ is the inverse of $I_t^*(\hat{\mu}_t, h^{t-1})$ over $\hat{\mu}_t$.

Ignoring some technical details that are dealt with in the Appendix (lemma 3), we find that a fully separating equilibrium is equivalent to the satisfaction of the manager’s first-order condition in investment, which in turn is given by the following differential equation:

$$[\hat{\mu}_t - I_t^*(\hat{\mu}_t, h^{t-1})] \frac{\partial I_t^*(\hat{\mu}_t, h^{t-1})}{\partial \hat{\mu}_t} = \lambda \frac{\partial E[\sigma | \hat{\mu}_t, h^{t-1}]}{\partial \hat{\mu}_t}, \quad \forall \hat{\mu}_t \in \mathbb{R}, \quad (4)$$

where $I_t^*(\hat{\mu}_{t-1}, h^{t-1}) = \hat{\mu}_{t-1}$ and $h^{t-1} = \{\hat{\mu}_{t-1}, \dots, \hat{\mu}_0\}$. Note that if no inferences are drawn on ability from investments, the first-best arises since the right-hand side of (4) is zero. However, from lemma 2, this is not generally the case, and the type of distortion depends on the sign of $\partial E[\sigma | \hat{\mu}_t, h^{t-1}] / \partial \hat{\mu}_t$. As will become clear below, depending on the period and the relative variances, σ^2 / τ^2 , the manager may prefer to signal either a “conservative” (i.e., $|\hat{\mu}_t - \hat{\mu}_{t-1}|$ small) or an “exaggerated” (i.e., $|\hat{\mu}_t - \hat{\mu}_{t-1}|$ large) posterior.

III. Overreaction and Exaggeration

Since the market does not observe the manager’s information itself, it uses investment, the outcome of the information, to update its opinion about managerial quality. Suppose that at period t the relative variance is such that $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma) / d\sigma < 0$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$, so more able managers change more on average. Such a situation occurs, for example, at time $t = 1$; see (2). Then the manager prefers the market to believe that his posterior is far from the period t prior. He accomplishes this by exaggerating his investment decisions. To the extent that the chosen investment level I_t^* satisfies either $I_t^* > \hat{\mu}_t > \hat{\mu}_{t-1}$ or $I_t^* < \hat{\mu}_t < \hat{\mu}_{t-1}$, we say that the manager is *overreacting to information* and *exaggerating profitability*. Along these lines, we present the following result.

PROPOSITION 1. Suppose that the market’s inferences h^{t-1} are correct and, during period t , $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma) / d\sigma < 0$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. Then there exists a unique equilibrium outcome in period t in which the manager overreacts to his new information whenever $\hat{\mu}_t \neq \hat{\mu}_{t-1}$ by

exaggerating the profitability of the project:

$$I_t^*(\hat{\mu}_t, h^{t-1}) > \hat{\mu}_t > \hat{\mu}_{t-1} \quad \text{or} \quad I_t^*(\hat{\mu}_t, h^{t-1}) < \hat{\mu}_t < \hat{\mu}_{t-1}.$$

Because (2) implies that the conditions of proposition 1 are met for any relative variance, σ^2/τ^2 in period 1, we have an immediate corollary.

COROLLARY 1. In the first period, the manager always overreacts to his own information.

The intuition for this result is straightforward. The dispersion of investment is increasing in ability if the manager acts efficiently. Therefore, the manager exaggerates whatever information he has in his investment choices, scaling back investment too much if the information is unfavorable and investing too much if the efficient investment exceeds the prior.

A useful illustration of the manager's incentives can be seen from the following numerical example in which $t = 1$, $\lambda = 1$, and $\tau = 1$ and σ is uniformly distributed on $[1, 3]$. The equilibrium choice of net investment as a function of $\hat{\mu}_1$ is given in figure 1.

Exaggeration implies that the manager overresponds to new information. This can be seen from figure 1, where I^* , the change in investment from the ex ante optimal level of zero, is higher in absolute value than $\hat{\mu}_1$, the first-best level. Note that the incentive to distort is greatest for intermediate changes; this arises because (i) for $\hat{\mu}_1$ close to $\hat{\mu}_0 = 0$, the manager's first-best investment level implies such slow learning that there is little that can be done (locally) to change opinions; and (ii) for $\hat{\mu}_1$ far from zero, there is little need to distort since the manager's first-best investment level already reveals fast learning.⁶

IV. Conservatism and Inertia

We now turn to the opposite environment in which $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)/d\sigma > 0$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. In such a setting, the manager would prefer the market to believe that his posterior is close to the period's prior; that is, $|\hat{\mu}_t - \hat{\mu}_{t-1}|$ is small. To this end, all managers will shade their investment decisions toward the market's prior belief of the efficient investment level. Using the first-order condition in (4), we can state the following characterization of equilibrium investments.

PROPOSITION 2. Suppose that the market's inferences h^{t-1} are correct; λ is sufficiently small that $\sup_{\hat{\mu}_t} \lambda(\partial^2/\partial \hat{\mu}_t^2)E[\sigma|\hat{\mu}_t, h^{t-1}] < 1/4$; and during period t , $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)/d\sigma > 0$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. Then there exists

⁶ More formally, this occurs because the density of the normal distribution is relatively flat at both the origin and the tails.

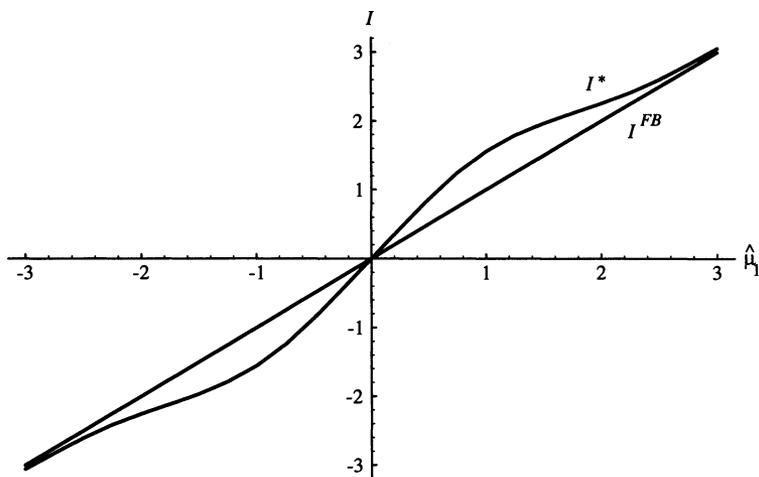


FIG. 1.—Exaggeration of information

a unique equilibrium outcome in period t . In this outcome, the manager is conservative and relies too little on his new information:

$$\hat{\mu}_t > I_t^*(\hat{\mu}_t, h^{t-1}) > \hat{\mu}_{t-1} \quad \text{or} \quad \hat{\mu}_t < I_t^*(\hat{\mu}_t, h^{t-1}) < \hat{\mu}_{t-1}.$$

Propositions 1 and 2 illustrate the inefficiencies in the model. Although corollary 1 illustrates that when managers initially choose projects they have a tendency to “go out on a limb,” proposition 2 indicates that they fail to admit mistakes by “staying put” if the variance of the posterior is declining in ability. The manager generally becomes conservative, not because of concavity in the rewards function, but rather because he “should know what he is doing” and would not change his opinion much if he were acting honestly.

There are two conditions for proposition 2 to necessarily hold. First, the market must infer h^{t-1} , the manager’s previous beliefs. This requires that there is a strictly monotonic relationship between investment in period $t - 1$ and $\hat{\mu}_{t-1}$; in other words, separation is required in period $t - 1$. If there is pooling in period $t - 1$, then it is impossible to precisely anchor the prior to determine how much the manager changes his opinion in period t .⁷ The second condition for the propo-

⁷ The implication of this is that we now must consider probability distributions over perceived changes. For example, suppose that all managers who observe $\hat{\mu}_{t-1} - \hat{\mu}_{t-2}$ between minus one and one do not change their investments in period $t - 1$. Then further suppose that the period t investment increases by $x > 0$ from period $t - 1$. The problem this gives rise to is that the market does not know where the manager lay on the range of previous-period changes from minus one to one, and so it cannot tell how much the manager has changed his beliefs in this period. Not only does this

sition is that λ cannot be too large. Proposition 2 provides the upper bound on λ such that separation will occur in period t as the unique outcome.⁸

Once again, we illustrate the distortions using an example in which $t = 2$, $\lambda = 1$, and $\tau = 3$ and σ is uniformly distributed on $[1, 3]$ as the market's prior distribution (see fig. 2). Here we observe behavior opposite to that in figure 1, with the movement in I^* being lower in absolute value than the warranted first-best investment I^{FB} , so that there is underreaction.

V. A Dynamic Model of Exaggeration and Conservatism

So far we have shown that exaggeration always occurs in period 1 and that conservatism can occur if the variance of the posterior is decreasing in managerial ability. In this section we (i) identify those factors that make conservatism or exaggeration likely and (ii) show that after some period of time, conservatism must occur forevermore.

A. Combining Exaggeration and Conservatism in a Two-Period Model

To understand the factors leading to a switch from exaggeration to conservatism, consider the second period for which the following proposition is immediate from (2).

PROPOSITION 3. If $\underline{\sigma}^2 > \tau^2\sqrt{2}$, the manager places too much emphasis on his period 2 observation and exaggerates information. If $\bar{\sigma}^2 < \tau^2\sqrt{2}$, the manager becomes excessively conservative in period 2 and relies too little on his period 2 observation.

These conditions have a simple intuitive interpretation.⁹ If the measurement error σ^2 for all managers is sufficiently small, the outcome is conservatism; if it is sufficiently high, exaggeration arises. One can also frame this result in terms of the type of task carried out by the manager: simpler tasks could be characterized by faster learning, but more complex tasks would be characterized by slower

entail placing distributions over possible beliefs in period $t - 1$, but the observation of changes in period t provides information on those beliefs, rendering inferences far more complicated.

⁸ If λ becomes too large, then there is nonexistence of a separating equilibrium. Instead, pooling arises as in Bernheim (1994).

⁹ This result follows immediately from the conditions on the support of σ . When $\underline{\sigma}^2 > \tau^2\sqrt{2}$, by (2) the variance of the posterior is declining in σ , and hence exaggeration occurs by proposition 1. When $\bar{\sigma}^2 < \tau^2\sqrt{2}$, the variance of the posterior is increasing in σ , so conservatism occurs by proposition 2.

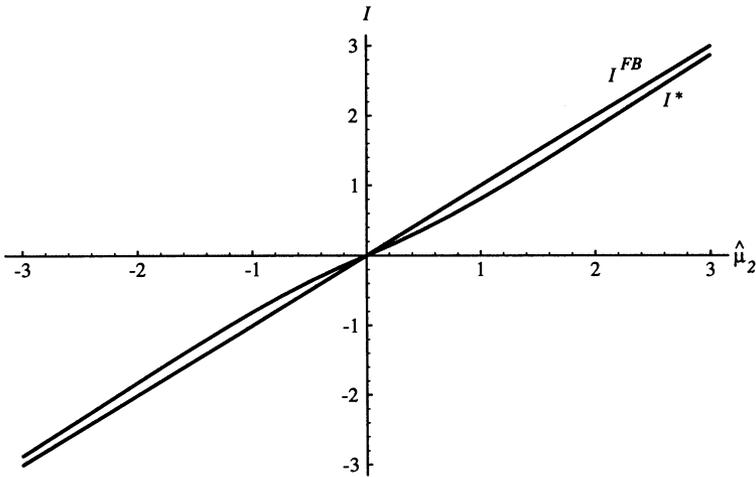


FIG. 2.—Ignored information and conservative investment

learning. Then proposition 3 shows that for simpler tasks, any deviation from period 1 choices illustrates managerial error. On the other hand, with complex tasks, there is no expectation that the manager should necessarily have been close the first time, so exaggeration occurs to suggest that the manager has “some more ideas of his own.”

B. A t-Period Model of Exaggeration and Conservatism

The characterization of the two-period model above is incomplete in that we have said nothing to describe period 2 behavior when $\underline{\sigma}^2 < \tau^2\sqrt{2} < \bar{\sigma}$, so that learning is neither very fast nor very slow. To more generally characterize outcomes, it is useful to define two variables. First, let \underline{t}^* be defined as the largest integer such that $\underline{\sigma}^2 > \tau^2\sqrt{\underline{t}^*(\underline{t}^* - 1)}$. This is the last time period for which a higher variance of investment necessarily implies more talent. Second, let \bar{t}^* be defined as the smallest integer such that $\bar{\sigma}^2 < \tau^2\sqrt{\bar{t}^*(\bar{t}^* - 1)}$. This characterizes the first time period beyond which more talented managers must show less variability in investment.

Our objective is then to determine the behavior of the manager in terms of these time “regions”: from period 1 to \underline{t}^* , from $\underline{t}^* + 1$ to $\bar{t}^* - 1$, and from \bar{t}^* onward. However, an additional problem arises within the possible “transition” period between $\underline{t}^* + 1$ and $\bar{t}^* - 1$, as it becomes uncertain whether the manager wishes to exaggerate or act conservatively since the sign of $d\bar{\sigma}_{\hat{\mu}_2}^2(\sigma)/d\sigma$ depends on σ . Furthermore, remember that our results on conservatism are predicated on the existence of separation in previous periods, so we need to

determine whether separation will occur in this transition region, where it is not clear whether an agent wishes to exaggerate or be conservative. However, under conditions similar to those described above, separation continues to hold, so that we can indeed claim that conservatism must occur after \bar{t}^* .

PROPOSITION 4. Suppose that λ is sufficiently small that $\sup_{\hat{\mu}_t} \lambda(\partial^2/\partial\hat{\mu}_t^2)E[\sigma|\hat{\mu}_t, \dots, \hat{\mu}_0] < 1/4$ for any $\{\hat{\mu}_t, \dots, \hat{\mu}_0\}$. Then there is a unique separating equilibrium such that, for all $t \leq \underline{t}^*$, exaggeration occurs and, for all $t \geq \bar{t}^*$, conservatism occurs, where $1 \leq \underline{t}^* < \bar{t}^*$.

This proposition segments the manager's behavior over time into three regions. Initially he exaggerates his own information since changing behavior shows that he has some ideas of his own. This occurs from period 1 to period \underline{t}^* . Following this, the manager may enter a transition period in which we have not characterized his behavior except to illustrate the conditions under which separation arises. However, after period \bar{t}^* , he becomes conservative forevermore. As a consequence, the model implies that ultimately managers become conservative. The key to understanding this result is that, over time, the prior entering each successive period is increasingly associated with the contribution of the manager. In period 1, he has made no contribution, but eventually his own opinion becomes dominant. Thus over time he is increasingly held to blame for "changing his mind." It is also worth noting here that the nature of the distortion changes over time; it is not simply that the size of the distortion changes with more observations, as is common in learning models.¹⁰

*C. Determinants of Exaggeration and Conservatism:
Comparative Statics*

It is important to address those factors that lend themselves to one form of distortion or the other. This issue is addressed in the following proposition.¹¹

¹⁰ It is important to point out that the critical times concern tenure on a project rather than managerial age per se. The conservatism that arises here is that when a manager has been on a project for a long time, he becomes reluctant to admit mistakes. In our model, experienced managers who are moved to new projects will once again act in an exaggerated fashion, so the model should truly be seen as addressing the effects of tenure on a project rather than aging. However, it is worth pointing out that our results rely on uncertainty about the ability of the manager. But $(\partial/\partial\hat{\mu}_t)E[\sigma|\hat{\mu}_t, \dots, \hat{\mu}_0] \rightarrow 0$ as $t \rightarrow \infty$, so the distortionary effects are reduced for older managers, with project tenure held constant. As a consequence, our model suggests that the inefficiencies described above are muted for older managers.

¹¹ Because \underline{t}^* and \bar{t}^* are integers, the comparative statics are stated in terms of weak monotonicity. Generically, for local changes, the critical times are unaffected. For sufficiently large changes of the parameters under study, the critical times will jump discontinuously. We are interested in the direction of these discontinuous jumps.

PROPOSITION 5. The critical times $\underline{t}^*(\underline{\sigma}, \tau)$ and $\bar{t}^*(\bar{\sigma}, \tau)$ are weakly increasing in the support of σ and weakly decreasing in τ .

The results of proposition 5 follow directly from differentiation and illustrate the conflict between the quality of public and private information that generates our results. The value of the public information in our model is given by the prior, with variance τ^2 , and the private information is parameterized by σ^2 . The more important private information is, the longer the region of conservatism will be. On the other hand, valuable public information leads to exaggeration. Any increase in the prior variance results in a shorter regime in which exaggeration occurs and a longer regime in which conservatism occurs as the manager's contribution builds up quickly. By contrast, the opposite occurs when the importance of private information falls, that is, as σ rises. In this sense an improvement in the quality of public information (private information) leads to a longer period of exaggeration (conservatism).¹²

In addition, our model also predicts that repetition of the same action is likely to give rise to conservatism. Furthermore, both types of distortions are likely to be most important when our actions are observed by others, who are making inferences on our capabilities. Empirical work by Kiesler (1971), Salancik (1977), and Berg, Dickhaut, and Kanodia (1991) clearly shows that the problem of sunk costs becomes more pervasive when (i) the act is public and visible to others and (ii) the act has been performed a number of times. We see both of these results as supportive of our work.

VI. Extensions

The previous section yielded some simple comparative statics showing how each type of distortion depended on the speed of learning. The purpose of this section is to yield further insights into how decisions are made. We do this by considering two plausible extensions of the model: (i) allowing the environment to change and (ii) allowing for information that is not privately collected by the manager, such as

¹² To better understand this conflict between the quality of public and private information, consider a limiting case in which the public information becomes unimportant, i.e., in which $\tau^2 \rightarrow \infty$. In this case the only valuable information is the manager's, which should bias our results toward conservatism. To see this, note that in period 1, $(\partial/\partial\hat{\mu}_1)E[\sigma|\hat{\mu}_1, \hat{\mu}_0] \rightarrow 0$ as $\tau^2 \rightarrow \infty$ because with a disparate prior there is almost no information on talent from the manager's choice of first-period investment. If this is the case, then exaggeration disappears and the manager chooses (almost) the efficient level of investment in period 1. But note that as $\tau^2 \rightarrow \infty$, proposition 2 implies that conservatism always occurs in period 2 and all future periods. Therefore, eliminating public information gives rise to only the exercise of conservatism.

on lagged revenues. Each subsection allows us to further refine our understanding of the likelihood of each type of distortion.

A. *Changing Environments*

In many instances, managers make decisions in changing environments, and the correct decision yesterday may not be appropriate today. In this subsection, we consider how a changing environment affects the distortionary behavior taken by the manager by allowing investment returns in period t (which in the previous sections were simply μ) to move according to a random walk $\mu_t = \mu_{t-1} + \zeta$, where $\zeta \sim \mathcal{N}(0, z^2)$. A useful parameterization of z^2 relates it to the market's prior variance, so we denote $z^2 \equiv \alpha\tau^2$. Note that when $\alpha = 0$, we are in our standard model with an unchanging environment. An issue that then arises is that although the manager may be relatively sure that he made the right decision last period, his current observation may still be of considerable importance since profitability may have changed.

We illustrate this issue in a two-period framework. We adapt the model of the previous section by assuming that the profits are now given by $\mu_t I_t - \frac{1}{2} I_t^2$, where $\mu_1 = \mu$ and $\mu_2 = \mu + \zeta$. The manager's posterior in period 1, $\hat{\mu}_1$, is therefore a less precise measure of profitability for period 2 than in previous sections. In addition, the manager draws an observation in period 2, $m_2 = \mu_2 + \epsilon_2$, where true profitability is observed with noise in period 2. Therefore, the manager knows that the world is evolving but still realizes that his previous estimate has value.

PROPOSITION 6. In a two-period setting in which the economic environment follows a random walk with variance $\alpha\tau^2$, there exists a decreasing function $\kappa(\alpha)$ with $\kappa(0) = \sqrt{2}$ such that, if $\underline{\sigma}^2 > \tau^2\kappa(\alpha)$, the manager exaggerates in period 2, and if $\bar{\sigma}^2 < \tau^2\kappa(\alpha)$, the manager is conservative in period 2. Furthermore, for $\alpha \geq 1$, the manager exaggerates in period 2 independent of σ .

The effect of the changing environment is characterized by the κ function. When $\alpha = 0$, profitability does not evolve over time, so that our results on period 2 behavior are unchanged; that is, if $\bar{\sigma}^2 > \tau^2\sqrt{2}$, exaggeration will always occur. However, as α rises, the environment becomes more unpredictable, so that the region of exaggeration occurs for a wider range of parameter values, $\bar{\sigma}^2 > \tau^2\kappa(\alpha)$, where $\kappa(\alpha) < \sqrt{2}$. If the innovation in the environment is as great as the initial prior (i.e., $\alpha \geq 1$), the manager exaggerates in both periods regardless of σ . Therefore, a slowly evolving environment is critical for the existence of conservatism. The reason for this is that when the underlying profitability of the projects evolves, the manager is held less to blame

for changing investment. Indeed, if the environment evolves enough, he will be held to blame if he does not change investment from one period to the next, even if his inferences are extremely good in any period. Therefore, in a rapidly evolving environment, managers show talent by always changing behavior more than is optimal.

B. Observable Firm Profitability

So far, we have assumed that the manager receives only his own reads on the profitability of the project. As a result, the investments made over time become increasingly associated with the previous contributions of the manager. However, in reality we might imagine that other information on profitability might arise (such as lagged observations of profits). In this subsection, we consider the implication of noisy information on the profitability of the project, where the quality of the information does not reflect managerial talent. Two additional insights arise from this extension. First, with respect to this new information, the manager always acts conservatively; he will exaggerate only with respect to his own privately collected information. Second, it is no longer the case that the manager must eventually become conservative with respect to his private information; instead, he may exaggerate his information indefinitely.

In order to model this phenomenon, we adapt our model in the previous section by assuming that in all odd-numbered periods the manager receives his private observations on profitability, m_t , as described in the previous sections; in even-numbered periods, the manager observes $\theta_t = \mu + \eta_t$, where η_t is normally distributed with mean zero and variance τ^2 . To retain simplicity, we assume that all information is privately held when he makes the investment decision. Note, however, that the quality of the θ_t information is not related to his talent. The simplest interpretation of this is that the manager receives noisy information on the revenues of the project, but does so before the market.¹³ Note here that the observations on profits can be used to make inferences about the manager's ability to learn, so that not all learning occurs through investment. However, given the preferences of the manager, learning from previous data (whether from output or investments) has an effect only through the prior on the manager's ability entering a given period. We do not

¹³ For example, assume that the market observes a noisy observation of θ_t at the beginning of period $t + 1$. Then our results are unchanged by allowing the market to also observe information on profits. Note also that it is important that the market does not observe exactly what the manager saw, because then the market can invert the manager's investment choice in period t to identify σ for sure.

explicitly characterize the prior on σ using all previous data for the reason that we identify the type of distortion in period t for *any* nondegenerate prior on σ . Therefore, our results below allow learning from revenue data.¹⁴

1. Response to Own Information

In order to maintain comparability with the previous sections of the paper, the appropriate period for comparison with period t in the previous sections is now period $2t - 1$, since there are t private observations obtained by that time. In order to understand the effect of t managerial observations, we consider period $2t - 1$, where previous information is incorporated in $\hat{\mu}_{2t-2}$. Then optimal updating yields the result that μ is distributed normally with mean

$$\hat{\mu}_{2t-1} = \left[\frac{\tau^2}{t(\tau^2 + \sigma^2)} \right] m_{2t-1} + \left[\frac{t\sigma^2 + (t-1)\tau^2}{t(\tau^2 + \sigma^2)} \right] \hat{\mu}_{2t-2}$$

and variance $V[\mu | \hat{\mu}_{2t-1}, \sigma] = \sigma^2 \tau^2 / t(\tau^2 + \sigma^2)$. Taking the variance of the posterior in the equation above and conditioning on $\hat{\mu}_{2t-2}$ and σ yields the conditional variance of $\hat{\mu}_{2t-1}$,

$$V[\hat{\mu}_{2t-1} | \hat{\mu}_{2t-2}, \sigma] = \frac{\sigma^2 \tau^4}{t[(t-1)\tau^2 + t\sigma^2](\tau^2 + \sigma^2)}.$$

Note that

$$\frac{\partial V[\hat{\mu}_{2t-1} | \hat{\mu}_{2t-2}, \sigma]}{\partial \sigma} > 0 \quad \text{iff} \quad \sigma^2 > \tau^2 \frac{\sqrt{t(t-1)}}{t}.$$

This condition implies exaggeration if it holds for all σ , that is, $\sigma^2 > \tau^2 \sqrt{t(t-1)}/t$. By comparison, when the manager does not observe the data on revenues, the relevant condition is $\sigma^2 > \tau^2 \sqrt{t(t-1)}$, so that, for all $t > 1$, conservatism becomes less likely when public information is available. Furthermore, note that $\sqrt{t(t-1)}/t \rightarrow 1$ as $t \rightarrow \infty$. Therefore, if $\sigma^2 > \tau^2$, conservatism *can never occur*. Thus conservatism can occur only if the manager's information is better than the revenue data.¹⁵ If the observation on profits has higher precision than the manager's observation, he is increasingly less to blame for changing his mind over time, which implies that the manager exaggerates forever. In the case in which the manager's observa-

¹⁴ Note, however, that the magnitude of the inefficiencies will depend on the informativeness of revenues, since the extent of his distortionary behavior is reduced as more is known about him.

¹⁵ By contrast, note that $\sqrt{t(t-1)} \rightarrow \infty$ as $t \rightarrow \infty$, so that conservatism must occur in our basic model.

tion is better than the market's, the temporal behavior of the manager described above continues to hold since he becomes conservative after period \bar{t}_o^* , where $\bar{t}_o^* \equiv \tau^4/(\tau^4 - \bar{\sigma}^4) > 1$.

2. Response to Revenue Data

Now consider the response of the manager to the θ_t information on revenues. In period $2t$, optimal updating yields the result that μ is distributed normally with mean

$$\hat{\mu}_{2t} = \left[\frac{\sigma^2}{t\tau^2 + (t+1)\sigma^2} \right] \theta_{2t} + \left[\frac{t(\tau^2 + \sigma^2)}{t\tau^2 + (t+1)\sigma^2} \right] \hat{\mu}_{2t-1}.$$

Straightforward manipulations yield

$$\bar{\sigma}_{\hat{\mu}_t}^2(\sigma) \equiv V[\hat{\mu}_{2t} | \hat{\mu}_{2t-1}, \sigma] = \frac{\sigma^4 \tau^2}{t[t\tau^2 + (t+1)\sigma^2](\tau^2 + \sigma^2)}.$$

Differentiating this function with respect to σ yields an expression that is always negative, so that the more talented the manager the less he changes his behavior. This implies that the manager always acts in a conservative fashion with respect to information whose precision is unrelated to his own ability.¹⁶

Once again, the reason for this should be clear. The only reason for the manager to exaggerate in our model is to signal that the latest piece of information is valuable. But the quality of the θ_t information does not reflect his talent, and so the only effect that remains is that of conservatism. Our findings under an environment of such observable profit signals are summarized in the following proposition.

PROPOSITION 7. Suppose that profits are observed with noise in even periods and λ is sufficiently small (as in proposition 4). Then there exists a unique separating equilibrium. In even periods, if $\tau \leq \sigma$, exaggeration occurs; if $\tau > \sigma$, then there exists \bar{t}_o^* and \bar{t}_o^* such that exaggeration occurs for all $t < \bar{t}_o^*$ and conservatism occurs for all $t > \bar{t}_o^*$. In odd periods, the manager always acts conservatively.

VII. Conclusion

Economists typically think of decision making in terms of net present value calculations in which value is generally measured in terms of

¹⁶ Batchelor and Dua (1992) consider how forecasters of economic aggregates update their forecasts on the basis of information about both the forecasts of other experts and other observable information. They find that forecasters overweight their previous estimates relative to the optimal estimator, and correspondingly "in all cases their error is to give less weight to the consensus [forecast] than would be necessary to minimize expected errors" (p. 170). This is supportive of our results. See also Lamont (1995) for related empirical work.

wealth or consumption flows generated by a trade. Social psychologists argue that this methodology ignores many aspects of human psychology that systematically affect the way decisions are made. A central concern of psychologists is the concept of cognitive dissonance reduction, which derives from a tension that arises when an individual holds two cognitions that are psychologically inconsistent. The impact of this tension is that individuals make decisions and interpret information in order to justify to themselves previous decisions they have made.

Essentially, cognitive dissonance reduction considers individuals as rationalizing beings rather than rational beings. An important implication of attempts to rationalize previous behavior is the process of escalating commitment (or the sunk cost fallacy), where individuals commit more resources to a losing cause so as to justify or rationalize their previous behavior (Staw and Ross 1989).¹⁷ However, another part of the psychology literature illustrates how individuals may over-respond to new information when making decisions. For example, numerous studies have analyzed how investors overreact to new information provided by the stock market rather than appropriately incorporating information on the population distribution (see, e.g., De Bondt and Thaler 1987). Similarly, Kahneman and Tversky (1982) ask individuals to identify the occupation of an individual in an experimental setting and show that they ignore information on the population distribution when they make predictions. This decision-making bias is often referred to as the base rate fallacy.

Our analysis yields outcomes similar to the base rate fallacy and the sunk cost fallacy. We do not claim that individuals do not exhibit the effects of cognitive dissonance reduction since many of the psychology results are derived under carefully controlled circumstances. Instead, our objective is to show that when individuals care about their reputations, they are likely to exhibit behavior in response to the economic incentives that is observationally equivalent to the psychological evidence.

A recent literature in economics on herding also addresses how managers may take actions in order to appear talented. First, Scharfstein and Stein (1990) consider a model in which, by assumption, untalented managers hold more disparate priors than their more talented counterparts, which implies that managers shade their decisions toward their prior. Our paper generates the dynamic correlations between variance and posteriors from more basic Bayesian updating. Furthermore, our results suggest that the opposite behav-

¹⁷ The dissonance that arises here is that the perceptions "I make good decisions" and "my project is failing" are cognitively dissonant.

ior must occur initially since concentrating behavior around the prior in period 1 shows a lack of talent, so our results predict “anti-herding” initially. Only after the manager has contributed to the project can conservatism occur. Therefore, although the outcome in some cases is similar to that in the herding literature, our results arise only through the impact of the manager’s previous contribution. Second, a series of papers link managerial decision making to nonlinearities in rewards from outcomes.¹⁸ For instance, Zwiebel (1995) considers how the possibility of being fired in a relative performance evaluation context induces nonlinearities in payoffs; as a result, (i) those who are behind take (efficient) risks and (ii) those who are slightly ahead of the required threshold for being retained become too conservative in the riskiness of their investments. Therefore, nonlinearities in payoffs affect risk taking. The difference between our work and this line of research is that we do not require nonlinearities in payoffs to get such results.¹⁹ Instead inefficiencies occur in our model simply through the learning process.²⁰

Finally, Kanodia, Bushman, and Dickhaut (1989) and Boot (1992) show how conservatism can occur for reasons similar to our own in the second period of a two-period model. They consider a binary investment choice model and assume that the true profitability of their project is always revealed in the second period. Their structures do not allow the possibility of exaggeration. In these circumstances, they show that conservatism can occur. The closest analogy to our model would arise if we assumed that perfect information on profitability became available in period 2. Straightforward calculations available from the authors show that adding this assumption to our model necessarily implies conservatism in period 2 and exaggeration in period 1. Thus our more general structure allows precise results even when the truth is constrained to be always revealed quickly.

To conclude, perhaps the most important points of our model of

¹⁸ See Holmström and Ricart i Costa (1986) for early work on the effect of nonlinearities in payoffs on investment.

¹⁹ In addition, we do not interpret conservatism as an unwillingness to take risks: we simply interpret it as an unwillingness to change previous behavior.

²⁰ Also related to our work are the communication-based inefficiencies of herding studied by Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Hirshleifer and Welch (1994). These papers outline how herding is likely when individuals can transmit information only through their actions. An individual who carries out a particular action is unable to tell others that he may have been close to indifferent between two actions, so that individuals follow his lead when there may be little reason to do so. A key distinction between our work and that mentioned above is that our paper holds that conservatism occurs when the individual has carried out actions on the basis of his own information before, whereas the work outlined above holds that herding is more likely when others carry out the previous actions, so that the true underlying information is in the hands of others.

reputation over learning are, first, that both conservatism and exaggeration arise in predictable rational ways and, second, that they both arise from precisely the same reputation problem. Therefore, while some may argue that decision making under reputation formation is likely to lead to herding or conservatism, it is our contention that those reputation formation features that give rise to conservatism are also likely to cause exaggeration earlier in the manager's tenure on the project. Thus the two types of behavior are inextricably linked.

Appendix

Proofs of Results

Proof of Lemma 2

Let $\phi(x, z^2)$ be the normal density for a random variable x distributed $\mathcal{N}(0, z^2)$, where ϕ_i represents the partial derivative of ϕ with respect to its i th argument. Let the prior distribution of σ beginning period t be given by $f(\sigma|h^{t-1})$, where $h^{t-1} \equiv \{\hat{\mu}_{t-1}, \dots, \hat{\mu}_0\}$. With Bayes's rule, the conditional density of σ is

$$f(\sigma|\hat{\mu}_t, h^{t-1}) = \frac{f(\sigma|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma))}{\int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds}$$

Consider the effect of $\hat{\mu}_t$ on this conditional density of σ :

$$\frac{\partial f(\sigma|\hat{\mu}_t, h^{t-1})}{\partial \hat{\mu}_t} = \frac{f(\sigma|h^{t-1})\phi_1(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma)) \int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds}{\left[\int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds \right]^2} - \frac{f(\sigma|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma)) \int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi_1(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds}{\left[\int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds \right]^2}$$

Hence, the sign of the numerator determines the sign of the partial derivative. In particular,

$$\text{sign} \left[\frac{\partial f(\sigma|\hat{\mu}_t, h^{t-1})}{\partial \hat{\mu}_t} \right] = \text{sign} \left[\frac{\phi_{\hat{\mu}_t}(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma))}{\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(\sigma))} - \frac{\int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi_1(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds}{\int_{\underline{\sigma}}^{\bar{\sigma}} f(s|h^{t-1})\phi(\hat{\mu}_t - \hat{\mu}_{t-1}, \bar{\sigma}_{\hat{\mu}_t}^2(s))ds} \right]$$

Because the second part of the right-hand-side expression is constant, we can say that the partial derivative $\partial f(\sigma|\hat{\mu}_t, \dots, \hat{\mu}_0)/\partial \hat{\mu}_t$ changes sign once

over the support $[\underline{\sigma}, \bar{\sigma}]$ provided that the first term is monotonic in σ . (Unless the derivative is exactly equal to zero for every σ , it must change sign at least once because f is a density that integrates to unity.)

The derivative of the first term with respect to $\bar{\sigma}_{\hat{\mu}_t}$ is simply $2(\hat{\mu}_t - \hat{\mu}_{t-1})/\bar{\sigma}_{\hat{\mu}_t}^3$. Consider first the case in which $\hat{\mu}_t > \hat{\mu}_{t-1}$, so this derivative is positive. If $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)/d\sigma > 0$, then the ratio increases in σ . In such a case, an increase in $\hat{\mu}_t$ shifts weight in the density f from low σ to high σ , thereby raising the expectation $E[\sigma | \hat{\mu}_t, \dots, \hat{\mu}_0]$. The alternative cases in which either $\hat{\mu}_t < \hat{\mu}_{t-1}$ or $d\bar{\sigma}_{\hat{\mu}_t}^2(\sigma)/d\sigma < 0$ holds are similarly straightforward, resulting in lemma 2. Q.E.D.

Proof of Propositions 1 and 2

The following lemma is useful in proving propositions 1 and 2.

LEMMA 3. The functions $\{I_t^*(\hat{\mu}_t, h^{t-1}), \hat{\mu}_t^*(I_t, h^{t-1})\}_{t=1}^T$ constitute a perfect Bayesian-Nash separating equilibrium iff, for all t and h^{t-1} , (1) the functions $\{I_t^*(\hat{\mu}_t, h^{t-1}), \hat{\mu}_t^*(I_t, h^{t-1})\}_{t=1}^T$ are continuous, almost everywhere differentiable, and strictly increasing in their first arguments; (2) $I_t^*(\hat{\mu}_{t-1}, h^{t-1}) = \hat{\mu}_{t-1}$ and equation (4) holds:

$$[\hat{\mu}_t - I_t^*(\hat{\mu}_t, h^{t-1})] \frac{\partial I_t^*(\hat{\mu}_t, h^{t-1})}{\partial \hat{\mu}_t} = \lambda \frac{\partial E[\sigma | \hat{\mu}_t, h^{t-1}]}{\partial \hat{\mu}_t} \quad \forall \hat{\mu}_t \in \mathbb{R}$$

wherever $\partial I_t^*(\hat{\mu}_t, h^{t-1})/\partial \hat{\mu}_t$ exists; and (3) $\hat{\mu}_t^*(I_t, h^{t-1})$ is the inverse of $I_t^*(\hat{\mu}_t, h^{t-1})$ over $\hat{\mu}_t$.

Proof of Lemma 3

Necessity.—We prove that each of the three conditions is implied by a separating equilibrium. Consider first condition 1. A revealed preference argument demonstrates that the functions must be increasing in their first arguments. Suppose otherwise. Taking h^{t-1} as given, choose $\hat{\mu}_t$ and $\hat{\mu}'_t$ such that $\hat{\mu}_t > \hat{\mu}'_t$, and I_t and I'_t are the respective optimal choices of net investment. Revealed preference implies that

$$V_t(I_t, \hat{\mu}_t, h^{t-1}) \geq V_t(I'_t, \hat{\mu}_t, h^{t-1})$$

and

$$V_t(I'_t, \hat{\mu}'_t, h^{t-1}) \geq V_t(I_t, \hat{\mu}'_t, h^{t-1}).$$

Simplification and rearrangement yield

$$(\hat{\mu}_t - \hat{\mu}'_t)I_t \geq V_t(I_t, \hat{\mu}_t, h^{t-1}) - V_t(I'_t, \hat{\mu}'_t, h^{t-1}) \geq (\hat{\mu}_t - \hat{\mu}'_t)I'_t,$$

which further implies that $(\hat{\mu}_t - \hat{\mu}'_t)(I_t - I'_t) \geq 0$. Thus the functions are weakly increasing. Because the equilibrium is fully separating, this relationship must be strict. Monotonicity in turn implies that the functions are almost everywhere differentiable. To see that it is continuous in $\hat{\mu}_t$, suppose to the contrary that there were a discontinuity. In a fully separating equilibrium, the utility difference between being thought a manager with posterior $\hat{\mu}_t$ and

one with posterior $\hat{\mu}_t + d\hat{\mu}$ is of order $\mathcal{O}(d\hat{\mu})$. For $d\hat{\mu}$ sufficiently small, this is less than the fixed differences in investment that the discontinuity implies. Thus one manager would prefer to choose the allocation of the other, contradicting the hypothesis of separation.

Consider condition 2. The fact that there is an equilibrium implies that an optimal I_t exists for each period. Because the support of σ is compact, the upper bound on any reputation effect is bounded. Because π_t is strictly concave in investment and unbounded below, we know that the optimal I_t must be finite. Therefore, a necessary condition for I_t is given by

$$\frac{\partial V_t(I_t, \hat{\mu}_t, h^{t-1})}{\partial I_t} = 0$$

or, alternatively,

$$\hat{\mu}_t - I_t = \lambda \frac{\partial E[\sigma | \hat{\mu}_t^*(I_t, h^{t-1}), h^{t-1}]}{\partial \hat{\mu}_t} \frac{\partial \hat{\mu}_t^*(I_t, h^{t-1})}{\partial I_t}.$$

This is precisely the first-order differential equation given in condition 2. The initial condition that $I_t^*(\hat{\mu}_{t-1}, h^{t-1}) = \hat{\mu}_{t-1}$ when $h^{t-1} \equiv \{\hat{\mu}_{t-1}, \dots, \hat{\mu}_0\}$ is implied by lemma 2.

Condition 3 follows directly from the second requirement in the definition of a perfect Bayesian separation equilibrium.

Sufficiency.—We shall show that any I_t^* satisfying conditions 1 and 2 must be globally optimal. This implies that the first requirement in the definition of a separating equilibrium is satisfied. (The second requirement is equivalent to condition 3.)

Suppose to the contrary that conditions 1 and 2 are satisfied but we do not have an optimum. Thus there exists some I_t' such that

$$V_t(I_t', \hat{\mu}_t, h^{t-1}) > V_t(I_t^*(\hat{\mu}_t, h^{t-1}), \hat{\mu}_t, h^{t-1}).$$

Because $I_t^*(\hat{\mu}_t, h^{t-1})$ is one-to-one in \mathbb{R} , there exists a posterior $\hat{\mu}_t'$ such that $I_t^*(\hat{\mu}_t', h^{t-1}) = I_t'$. By continuity, we can rewrite the inequality as an integral:

$$\int_{\hat{\mu}_t'}^{\hat{\mu}_t} \frac{\partial}{\partial I_t} V_t(I_t^*(s, h^{t-1}), \hat{\mu}_t, h^{t-1}) \frac{\partial I_t^*(s, h^{t-1})}{\partial \hat{\mu}_t} ds < 0.$$

The first-order condition 2 implies that $(\partial/\partial I_t)V_t(I_t^*(\hat{\mu}_t, h^{t-1}), \hat{\mu}_t, h^{t-1}) = 0$ for all $\hat{\mu}_t \in \mathbb{R}$. Thus, combining terms, we get

$$\int_{\hat{\mu}_t'}^{\hat{\mu}_t} \left[\frac{\partial}{\partial I_t} V_t(I_t^*(s, h^{t-1}), \hat{\mu}_t, h^{t-1}) \frac{\partial I_t^*(s, h^{t-1})}{\partial \hat{\mu}_t} - \frac{\partial}{\partial I_t} V_t(I_t^*(s, h^{t-1}), s, h^{t-1}) \frac{\partial I_t^*(s, h^{t-1})}{\partial \hat{\mu}_t} \right] ds < 0.$$

Integrating once again, we get

$$\int_{\hat{\mu}_t'}^{\hat{\mu}_t} \int_s^{\hat{\mu}_t} \frac{\partial^2}{\partial I_t \partial \hat{\mu}_t} V_t(I_t^*(s, h^{t-1}), t, h^{t-1}) \frac{\partial I_t^*(s, h^{t-1})}{\partial \hat{\mu}_t} dt ds < 0.$$

But since $(\partial^2/\partial I_t \partial \hat{\mu}_t) V_t = 1$ and $\partial I_t^*(\hat{\mu}_t, h^{t-1})/\partial \hat{\mu}_t > 0$ (by condition 1), this inequality cannot be satisfied. Q.E.D.

Proof of Proposition 1

We take the history h^{t-1} as correct and suppress its notation in the arguments below focusing only on $\hat{\mu}_{t-1}$ explicitly. We then prove that separation occurs in period t , which satisfies the stated exaggeration conditions in the proposition.

Following lemma 3, we know that the set of equilibria is fully characterized by the set of increasing solutions to (4) that satisfy $I_t^*(\hat{\mu}_{t-1}) = \hat{\mu}_{t-1}$. We first show that there exists a unique solution to the differential equation that satisfies these conditions. We then characterize the solution.

We can rewrite (4) as

$$I_t^{*'}(\hat{\mu}_t) = \frac{\lambda(\partial/\partial \hat{\mu})E[\sigma|\hat{\mu}_t, h^{t-1}]}{\hat{\mu}_t - I_t^*(\hat{\mu}_t)}$$

We focus on the domain $[\hat{\mu}_{t-1}, \infty)$. The argument is analogous over the complementary domain. Over any interval $[a, b] \subset (\hat{\mu}_{t-1}, \infty)$, the differential equation is Lipschitz and admits a unique solution for a given value of $I_t^*(a)$. Finding a unique solution to the differential equation over $[\hat{\mu}_{t-1}, \infty)$ is problematic only at $\hat{\mu}_{t-1}$, where there is a singularity and the differential equation fails to be Lipschitzian. At this point, we may have multiple solutions. With L'Hospital's rule, the limit of the slope as $\hat{\mu}_t \rightarrow \hat{\mu}_{t-1}$ satisfies

$$\lim_{\hat{\mu}_t \rightarrow \hat{\mu}_{t-1}} I_t^{*'}(\hat{\mu}_t) = \lim_{\hat{\mu}_t \rightarrow \hat{\mu}_{t-1}} \frac{(\partial^2/\partial \hat{\mu}_t^2)\lambda E[\sigma|\hat{\mu}_t, h^{t-1}]}{1 - I_t^{*'}(\hat{\mu}_t)}$$

Rearranging and using the term $I_t^{*'}(\hat{\mu}_{t-1})$ for the limit, we find the following quadratic equation:

$$I_t^{*'}(\hat{\mu}_{t-1})^2 - I_t^{*'}(\hat{\mu}_{t-1}) + \gamma = 0,$$

where $\gamma \equiv (\partial^2/\partial \hat{\mu}_t^2)\lambda E[\sigma|\hat{\mu}_{t-1}, \hat{\mu}_{t-1}, \dots]$. There are two potential real roots to this equation:

$$I_t^{*'}(\hat{\mu}_{t-1}) = \frac{1 \pm \sqrt{1 - 4\gamma}}{2}$$

In the case in which $d\bar{\sigma}_{\hat{\mu}}^2(\sigma)/d\sigma < 0$, it is straightforward to demonstrate that $\gamma < 0$. Thus there are two real roots, one positive and one negative. We know that the negative root cannot be an equilibrium, so we focus only on the positive value. After simple algebraic manipulation, the positive root satisfies $I_t^{*'}(\hat{\mu}_{t-1}) > 1$. Finally, we need to show that the $I_t^{*'}(\hat{\mu}_t)$ associated with the positive-root solution stays positive over the domain $(\hat{\mu}_{t-1}, \infty)$. By lemma 2 we know that $\partial E[\sigma|\hat{\mu}_t, \hat{\mu}_{t-1}, \dots]/\partial \hat{\mu}_t < 0$. Hence, provided that $I_t^*(\hat{\mu}_t) > \hat{\mu}_t$, we know that $I_t^{*'}(\hat{\mu}_t) > 1$. Because the positive-root solution leaves the point $\hat{\mu}_{t-1} < \hat{\mu}_t$ with a slope greater than unity, $I_t^*(\hat{\mu}_t) > \hat{\mu}_t$ for all $\hat{\mu}_t > \hat{\mu}_{t-1}$. Thus there is a unique nondecreasing solution such that $I_t^*(\hat{\mu}_t) > \hat{\mu}_t > \hat{\mu}_{t-1}$ for all $\hat{\mu}_t > \hat{\mu}_{t-1}$. The case in which $\hat{\mu}_t < \hat{\mu}_{t-1}$ proceeds analogously. Q.E.D.

Proof of Proposition 2

As in the proof of proposition 1, we take the history h^{t-1} as correct and suppress its notation in the arguments below, focusing only on $\hat{\mu}_{t-1}$ explicitly. We then prove that separation occurs in period t that satisfies the stated conservatism conditions in the proposition.

Following lemma 3, we know that the set of equilibria is fully characterized by the set of increasing one-to-one solutions to (4) that pass through $I_t^*(\hat{\mu}_{t-1}) = \hat{\mu}_{t-1}$. We first show that there exists a unique solution to the differential equation that satisfies these conditions. We then characterize the solution.

We can rewrite (4) as

$$I_t^{*'}(\hat{\mu}_t) = F(I_t^*(\hat{\mu}_t), \hat{\mu}_t) \equiv \frac{\lambda(\partial/\partial\hat{\mu}_t)E[\sigma|\hat{\mu}_t, h^{t-1}]}{\hat{\mu}_t - I_t^*(\hat{\mu}_t)}.$$

Following the argument presented in the proof to proposition 1, we have a potential problem solving this differential equation at the point $\hat{\mu}_{t-1}$, where a singularity occurs. Taking limits and using L'Hospital's rule, we find the following quadratic equation:

$$I_t^{*'}(\hat{\mu}_{t-1})^2 - I_t^{*'}(\hat{\mu}_{t-1}) + \gamma = 0,$$

where $\gamma \equiv (\partial^2/\partial\hat{\mu}_t^2)\lambda E[\sigma|\hat{\mu}_{t-1}, \hat{\mu}_{t-1}, \dots]$. With our assumption on $d\hat{\sigma}_{\hat{\mu}_t}^2(\sigma)/d\sigma > 0$, it is straightforward to show that $\gamma > 0$; in combination with our assumption on λ , we have $1/4 > \gamma > 0$. Thus there are two real roots, both positive, which we denote by r_1 and r_2 ($r_2 > r_1$). Algebraic manipulation demonstrates that $r_1 \in (0, 1/2)$ and $r_2 \in (1/2, 1)$.

Normalize $\hat{\mu}_{t-1} = 0$ without loss of generality. Consider the positive orthant with the following two equations: $\alpha(\hat{\mu}_t) \equiv \hat{\mu}_t - 2(d/d\hat{\mu}_t)E[\sigma|\hat{\mu}_t]$ and $\beta(\hat{\mu}_t) \equiv \hat{\mu}_t$. Both equations pass through the origin, $\alpha(\hat{\mu}_t) < \beta(\hat{\mu}_t)$, for all $\hat{\mu}_t \in (0, \infty)$, and $\lim_{\hat{\mu}_t \rightarrow \infty} |\alpha(\hat{\mu}_t) - \beta(\hat{\mu}_t)| = 0$ because $\lim_{\hat{\mu}_t \rightarrow \infty} (d/d\hat{\mu}_t)E[\sigma|\hat{\mu}_t] = 0$. Additionally,

$$\alpha'(\hat{\mu}_t) = 1 - 2 \frac{\partial^2}{\partial\hat{\mu}_t^2} E[\sigma|\hat{\mu}_t, h^{t-1}] > F(\alpha(\hat{\mu}_t), \hat{\mu}_t) = 1/2,$$

$$\beta'(\hat{\mu}_t) = 1 < F(\beta(\hat{\mu}_t), \hat{\mu}_t) = \infty,$$

and $F_t(I, \hat{\mu}_t) > 0$. In the terminology of Hubbard and West (1991), the functions $\{\alpha(I), \beta(I)\}$ define an "anti-funnel." We can apply their theorem 1.4.5 to ascertain that there exists a unique solution to the differential equation that lies in the anti-funnel and has initial condition $I_t^*(0) = 0$. Because the solution lies below the 45-degree line but above $\alpha(\hat{\mu}_t)$ (which converges to the 45-degree line), it is strictly increasing and forms a one-to-one mapping over \mathbb{R}_+ . The candidate solution corresponding to the higher root, r_2 , begins in this anti-funnel and so represents this one-to-one solution, and hence is an equilibrium.

Consider the low-root solution. The candidate solution corresponding to r_1 must necessarily remain below the anti-funnel. We shall show that it (i) cannot converge to I as $I \rightarrow \infty$ and (ii) consequently cannot be part of an equilibrium. First, to show part i, because the anti-funnel provides an upper

bound on the r_1 solution, if $I_t^*(\hat{\mu}_t)$ converges to $\hat{\mu}_t$, it must do so more slowly than $(d/d\hat{\mu}_t)E[\sigma|\hat{\mu}_t]$ goes to zero:

$$|\hat{\mu}_t - I_t^*(\hat{\mu}_t)| > 2 \frac{d}{d\hat{\mu}_t} E[\sigma|\hat{\mu}_t].$$

With the differential equation, this implies that $I_t^{*'} < 1/2$ for all $\hat{\mu}_t$. Because it starts out below the 45-degree line, this solution can never converge to $\hat{\mu}_t$. Second, to show part ii, because $I_t^{*'}(\hat{\mu}_t) < 1/2$, the distance between the solution and the 45-degree line is unbounded. This area represents the lost productivity from inefficient equilibrium investment used to maintain a reputation. Because the total reputational loss from the manager's choice of an $I \neq I_t^*(\hat{\mu}_t)$ is bounded by $\lambda\Delta\sigma$, there exists a type $\hat{\mu}_t$ sufficiently large, where choosing according to the posited solution is not optimal. (Because the high-root solution to the differential equation is increasing and satisfies the initial condition, lemma 3 implies that the solution must not be one-to-one; i.e., it has a vertical asymptote.) Consequently, the low-root solution to the differential equation is not an equilibrium.

Thus we have a unique equilibrium (corresponding to the high root, r_2). This equilibrium lies within the constructed anti-funnel and therefore exhibits conservatism (i.e., it is below the 45-degree line). Thus $\hat{\mu}_t > I_t^*(\hat{\mu}_t) > \hat{\mu}_{t-1}$ for all $\hat{\mu}_t > \hat{\mu}_{t-1}$. The case in which $\hat{\mu}_t < \hat{\mu}_{t-1}$ proceeds analogously. Q.E.D.

Proof (Sketch) of Proposition 4

The proposition follows along the same lines of propositions 1 and 2 once we have ascertained that h^{t-1} is known by the market at the start of period t . Thus it is sufficient to show that a unique separating equilibrium exists during any period of the transition phase $t \in (t^*, \bar{t}^*)$.

Consider the first period of the transition phase so that h^{t-1} is correctly inferred. At $\hat{\mu}_{t-1}$, the function $(\partial/\partial\hat{\mu}_t)E[\sigma|\hat{\mu}_t, h^{t-1}] = 0$. If this is the only extremum point, then the model is still one of exaggeration and proposition 1 applies. Suppose instead that there are multiple extrema, and restrict attention to the positive orthant (with respect to $\hat{\mu}_{t-1}$). At each extremum as $\hat{\mu}_t$ is increased, the equilibrium investment function must pass through the 45-degree line, and so we have an additional set of initial conditions. For each region lying between two extrema, we have a region of either conservatism or exaggeration depending on whether or not $(\partial/\partial\hat{\mu}_t)E[\sigma|\hat{\mu}_t, h^{t-1}]$ is positive or negative. Thus either proposition 1 or 2 will apply, and full separation will result. Thus h^t will be correctly inferred at the start of period $t + 1$. By induction, separation occurs over the entire transition phase. Q.E.D.

Proof of Proposition 6

Consider the manager's optimal estimate of the value of μ_2 . He has two pieces of information. First, the manager has $\hat{\mu}_1$, which is an unbiased estimate of μ_2 . Second, the manager observes m_2 , which is also an unbiased estimate of

μ_2 with variance σ^2 . Together, the optimal estimate of $\hat{\mu}_2$ is given by

$$\hat{\mu}_2 = \left[\frac{\tau^2 \sigma^2 + z^2 (\tau^2 + \sigma^2)}{(z^2 + \sigma^2)(\tau^2 + \sigma^2) + \tau^2 \sigma^2} \right] m_2 + \left[\frac{\sigma^2 (\tau^2 + \sigma^2)}{(z^2 + \sigma^2)(\tau^2 + \sigma^2) + \tau^2 \sigma^2} \right] \hat{\mu}_1.$$

From the market's perspective at the beginning of period 2 (the market has inferred only $\hat{\mu}_1$), the conditional variance of m_2 is given by

$$V[m_2 | \hat{\mu}_1, \sigma] = \frac{\tau^2 \sigma^2}{\tau^2 + \sigma^2} + z^2 + \sigma^2.$$

Thus the market's conditional variance of $\hat{\mu}_2$ is given by

$$V[\hat{\mu}_2 | \hat{\mu}_1, \sigma] = \frac{[\tau^2 \sigma^2 + z^2 (\sigma^2 + \tau^2)]^2}{(\tau^2 + \sigma^2)[(z^2 + \sigma^2)(\tau^2 + \sigma^2) + \tau^2 \sigma^2]}.$$

We have

$$V[\hat{\mu}_2 | \hat{\mu}_1, \sigma] = \frac{[(1 + \alpha)\sigma^2\tau^2 + \alpha\tau^4]^2}{[(2 + \alpha)\tau^2\sigma^2 + \sigma^4 + \alpha\tau^4](\tau^2 + \sigma^2)}.$$

In order to identify whether the manager will act conservatively or exaggerate, we differentiate $V[\hat{\mu}_2 | \hat{\mu}_1, \sigma]$ with respect to σ as before to determine the regions of exaggeration and conservatism. Differentiating with respect to σ yields

$$\frac{\partial}{\partial \sigma} V[\hat{\mu}_2 | \hat{\mu}_1, \sigma] = \frac{-2\tau^4\sigma^3[\sigma^2(1 + \alpha) + \alpha\tau^2][\sigma^4(1 + \alpha) + 3\alpha\tau^2\sigma^2 - 2(1 - \alpha)\tau^4]}{(\sigma^2 + \tau^2)^2[\sigma^4 + (2 + \alpha)\sigma^2\tau^2 + \alpha\tau^4]^2}.$$

The sign of this expression is determined by the sign of

$$-(1 + \alpha)\xi^2 - 3\alpha\xi + 2(1 - \alpha),$$

where $\xi \equiv \sigma^2/\tau^2$. This quadratic equation has one positive real root, given by

$$\kappa(\alpha) = \frac{\sqrt{8 + \alpha^2 - 3\alpha}}{2(1 + \alpha)},$$

which satisfies $\kappa(0) = \sqrt{2}$, $\kappa(1) = 0$, and $\kappa'(\alpha) < 0$. Thus proposition 6 follows from applying propositions 1 and 2, using lemma 2 and the constructed $\kappa(\alpha)$ above. Q.E.D.

References

Banerjee, Abhijit V. "A Simple Model of Herd Behavior." *Q.J.E.* 107 (August 1992): 797-817.
 Batchelor, Roy, and Dua, Pami. "Conservatism and Consensus-Seeking among Economic Forecasters." *J. Forecasting* 11 (February 1992): 169-81.
 Berg, Joyce; Dickhaut, John; and Kanodia, Chandra. "The Role of Private Information in the Sunk-Cost Phenomenon." Manuscript. St. Louis: Washington Univ., 1991.
 Bernheim, B. Douglas. "A Theory of Conformity." *J.P.E.* 102 (October 1994): 841-77.

- Bikhchandani, Sushil; Hirshleifer, David; and Welch, Ivo. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *J.P.E.* 100 (October 1992): 992-1026.
- Boot, Arnoud W. A. "Why Hang on to Losers? Divestitures and Takeovers." *J. Finance* 47 (September 1992): 1401-23.
- De Bondt, Werner F. M., and Thaler, Richard H. "Further Evidence on Investor Overreaction and Stock Market Seasonality." *J. Finance* 42 (July 1987): 557-81.
- DeGroot, Morris H. *Optimal Statistical Decisions*. New York: McGraw-Hill, 1970.
- Gibbons, Robert, and Murphy, Kevin J. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *J.P.E.* 100 (June 1992): 468-505.
- Harris, Milton, and Holmström, Bengt. "A Theory of Wage Dynamics." *Rev. Econ. Studies* 49 (July 1982): 315-33.
- Hirshleifer, David, and Welch, Ivo. "Institutional Memory, Inertia, and Impulsiveness." Working Paper no. 2-94. Los Angeles: Univ. California, 1994.
- Holmström, Bengt. "Managerial Incentive Problems: A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki: Swedish School Econ., 1982.
- Holmström, Bengt, and Milgrom, Paul R. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55 (March 1987): 303-28.
- Holmström, Bengt, and Ricart i Costa, Joan. "Managerial Incentives and Capital Management." *Q.J.E.* 101 (November 1986): 835-60.
- Hubbard, John H., and West, Beverly H. *Differential Equations: A Dynamical Systems Approach*. Pt. 1. *Ordinary Differential Equations*. New York: Springer-Verlag, 1991.
- Kahneman, Daniel, and Tversky, Amos. "Intuitive Prediction: Biases and Corrective Procedures." In *Judgment under Uncertainty: Heuristics and Biases*, edited by Daniel Kahneman, Paul Slovic, and Amos Tversky. Cambridge: Cambridge Univ. Press, 1982.
- Kanodia, Chandra; Bushman, Robert; and Dickhaut, John. "Escalation Errors and the Sunk Cost Effect: An Explanation Based on Reputation and Information Asymmetries." *J. Accounting Res.* 27 (Spring 1989): 59-77.
- Kiesler, Charles A. *The Psychology of Commitment: Experiments Linking Behavior to Belief*. New York: Academic Press, 1971.
- Lamont, Owen. "Macroeconomic Forecasts and Microeconomic Forecasters." Working Paper no. 5284. Cambridge, Mass.: NBER, October 1995.
- Salancik, Gerald R. "Commitment and the Control of Organizational Behavior and Belief." In *New Directions in Organizational Behavior*, edited by Barry M. Staw and Gerald R. Salancik. Malabar, Fla.: Kreiger, 1977.
- Scharfstein, David S., and Stein, Jeremy C. "Herd Behavior and Investment." *A.E.R.* 80 (June 1990): 465-79.
- Staw, Barry M., and Ross, Jerry. "Understanding Behavior in Escalation Situations." *Science* 246 (October 13, 1989): 216-20.
- Zwiebel, Jeffrey. "Corporate Conservatism and Relative Compensation." *J.P.E.* 103 (February 1995): 1-25.