Learning To Be Overconfident*

Simon Gervais

University of Pennsylvania

and

Terrance Odean

University of California at Davis

First Version: 10 May 1996

This Version: 7 June 1999

*The authors would like to thank Roger Edelen, Hayne Leland, Andrew Lo, Ananth Madhavan, David Romer, Matthew Spiegel, Brett Trueman, Greg Willard, and two anonymous referees, as well as seminar participants at l'Ecole des Hautes Etudes Commerciales (Montréal), Rutgers University, the University of Texas at Austin, the 1998 meetings of the American Finance Association, and the UCLA conference on the "Market Efficiency Debate" for their comments and suggestions. All remaining errors are the authors' responsibility. Address correspondence to Simon Gervais, Finance Department, Wharton School, University of Pennsylvania, Steinberg Hall - Dietrich Hall, Suite 2300, Philadelphia, PA 19104-6367, (215) 898-2370.

Abstract

We develop a multi-period market model describing both the process by which traders learn about their ability, and how a bias in this learning can create overconfident traders. A trader in our model initially does not know his own ability. He infers this ability from his successes and failures. In assessing his ability the trader takes too much credit for his successes, and this leads him to become overconfident. A trader's expected level of overconfidence increases in the early stages of his career. Then, with more experience, he comes to better recognize his own ability. The patterns in trading volume, expected profits, price volatility, and expected prices resulting from this endogenous overconfidence are analyzed.

Your herds and flocks may increase, and you may amass much silver and gold — everything you own may increase. But your heart may then grow haughty, and you may... say to yourself, 'It was my own strength and personal power that brought me all this prosperity.'

Deuteronomy 8:13-14,17
The Living Torah
translation by Rabbi Aryeh Kaplan

An old Wall Street adage advises "Don't confuse brains with a bull market." The need for such wisdom stems from traders' willingness to attribute too much of their success to their own abilities and not enough to their good fortune. Successful traders are, thus, prone to become overconfident in their abilities.

It is a common feature of human existence that we constantly learn about our own abilities by observing the consequences of our actions. For most people there is an attribution bias to this learning: we tend to overestimate the degree to which we are responsible for our own successes (Wolosin, Sherman and Till, 1973; Miller and Ross, 1975; Langer and Roth, 1975). As Hastorf, Schneider and Polifka (1970) write, "we are prone to attribute success to our own dispositions and failure to external forces." People recall their successes more easily than their failures and "even misremember their own predictions so as to exaggerate in hindsight what they knew in foresight" (Fischhof, 1982).

In this paper, we develop a multi-period market model describing both the process by which traders learn about their ability and how a bias in this learning can create overconfident traders. Traders in our model initially do not know their ability. They learn about their ability through experience. Traders who successfully forecast next period dividends improperly update their beliefs; they weight too heavily the possibility that their success was due to superior ability. In so doing they become overconfident.

In our model, a trader's level of overconfidence changes dynamically with his successes and failures. A trader is not overconfident when he begins to trade. Ex ante, his expected overconfidence increases over his first several trading periods and then declines.[1] Thus the greatest expected overconfidence in a trader's lifespan comes early in his career. After this he tends to develop a progressively more realistic assessment of his abilities as he ages.

[1]More precisely, this will happen for insiders with learning biases that are not too large, where "not too large" is precisely defined in section 3.

One criticism of models of non-rational behavior is that non-rational traders will underperform rational traders and eventually be driven to the margins of markets if not out of them altogether.[2] This is, however, not always the case. De Long, Schleifer, Summers and Waldmann (1990) present a model where non-rational traders in an overlapping generations model earn higher expected profits than rational traders by bearing a disproportionate amount of the risk that they themselves create. Rational traders are unwilling to take long-term arbitrage positions to eliminate these higher profits because of the risk that they may die before the arbitrage pays off.

In our model, the most overconfident and non-rational traders are not the poorest traders. In fact, for any given level of learning bias and trading experience, it is successful traders, though not necessarily the most successful traders, who are the most overconfident. Overconfidence does not make traders wealthy, but the process of becoming wealthy *can* make traders overconfident.

The model also shows that success affects traders' conditional future expected profits in two ways. First, success is indicative of higher ability and, therefore, greater expected future profits. Second, success can increase overconfidence and thereby lower expected future profits through suboptimal decision-making. The detrimental effect of the more successful traders' greater overconfidence on their future expected profits may, on occasion, more than offset their greater probable ability.

Most models of financial market microstructure assume that all trader characteristics are common knowledge; in particular, traders' risk aversion, their wealth, and the distribution of their information are known by all market participants. Exceptions include Blume, Easley and O'Hara (1994), Gervais (1996), and Subramanyam (1996). In these papers, the precision of the traders' information is random. Each traders' precision is known to himself but is uncertain to other market participants who must infer it from his actions. Our model builds on these works by extending this uncertainty to the trader himself. He initially does not know the precision of his own information, and must infer it by observing his signals and subsequent outcomes.

A large literature demonstrates that people are usually overconfident and that, in particular, they are overconfident about the precision of their knowledge.[3] Benos (1998), Odean (1998), Kyle and Wang (1997), and Wang (1997) examine models with statically overconfident traders. In these models greater overconfidence leads to greater expected trading volume and greater price volatility.[4]

[2] Early proponents of this view include Alchian (1950), and Friedman (1953). More recently, Blume and Easley (1982, 1992) have reinforced these ideas analytically.

[3] See for example Alpert and Raiffa (1982), and Lichtenstien, Fischhoff and Phillips (1982). Odean (1998) provides an overview of this literature.

[4] In one exception, Odean (1998) shows that an overconfident, risk-averse market-maker may reduce market volatility.

In our model, a greater learning bias causes greater expected overconfidence, which leads to greater expected trading volume and greater price volatility.

Daniel, Hirshleifer and Subrahmanyam (1998) look at trader overconfidence in a dynamic model. Our paper differs from theirs in that we concentrate on the dynamics by which self-serving attribution bias engenders overconfidence in traders, and not on the joint distribution of trader ability and the risky security's final payoff.[5] Our approach has the advantage of being analytically tractable. Our analysis differs from that of Daniel, Hirshleifer and Subrahmanyam in that we are interested in the dynamic effects of biased learning on overconfidence, trading volume, price volatility and trading profits, whereas their main objective is to show that overconfidence implies different price correlation patterns in the short run and the long run. Nevertheless, we show that expected prices exhibit the same humped-shape patterns through time in both models.

Overconfidence is determined, in our model, endogenously and changes dynamically over a trader's life. This enables us to make predictions about when a trader is most likely to be overconfident (when he is inexperienced and successful) and how overconfidence will change during a trader's life (it will, on average, rise early in a trader's career and then gradually fall). The model also has implications for changing market conditions. For example, most equity market participants have long positions and benefit from upward price movements. We would therefore expect aggregate overconfidence to be higher after market gains and lower after market losses. Since, as we show, greater overconfidence leads to greater trading volume, this suggests that trading volume will be greater after market gains and lower after market losses. Indeed, Statman and Thorley (1998) find that this is the case. We would expect aggregate overconfidence to be particularly high in a market with many young traders who have experienced only a long bull run. Thus, our model predicts that trading volume and volatility should be higher in the late stages of a bull market than in the late stages of a bear market.

Rabin and Schrag (1999) develop a model of confirmatory bias, the tendency to interpret new information as confirming one's previous beliefs. In as much as people tend to have positive self-images, confirming bias and self-serving attribution bias are related. Our paper differs considerably from Rabin and Schrag's, in that we analyze the effect of attribution bias on the overconfidence of traders in financial markets.

The rest of this paper is organized as follows. In section 1, we introduce a one-security multi-period economy with one insider, one liquidity trader and one market maker. In section 2, we develop the conditions under which there is a unique linear equilibrium in our economy. This linear equilibrium is used in section 3 to analyze the effects of the insider's learning bias on his

[5]This is because the risky security's dividend is publicly revealed at the end of every trading round in our model.

overconfidence and profits, as well as on the market's trading volume, volatility, and price patterns. Section 4 discusses the empirical implications of the model and, finally, section 5 concludes. All the proofs are contained in the appendices.

1 The Economy

We study a multi-period economy in which only one risky asset is traded among three market participants: an informed trader, a liquidity trader, and a market maker. At the end of period t, the risky asset pays off a dividend $\hat{v}_t$, unknown to all the market participants at the beginning of the period.[6]

At the beginning of each period t, the risk-neutral informed trader (also called *insider*) observes a signal $\hat{\theta}_t$ which is correlated with $\hat{v}_t$. In particular, the signal $\hat{\theta}_t$ is given by

$$\hat{\theta}_t = \hat{\delta}_t \hat{v}_t + (1 - \hat{\delta}_t)\hat{\varepsilon}_t, \tag{1}$$

where $\hat{\varepsilon}_t$ has the same distribution as $\hat{v}_t$, but is independent from it. The variable $\hat{\delta}_t$ takes the values 0 or 1. Since $\hat{\varepsilon}_t$ is independent from $\hat{v}_t$, the insider's information will only be useful when $\hat{\delta}_t$ is equal to one. We assume that this will happen with probability $\hat{a}$, i.e.

$$\hat{\delta}_t \mid \hat{a} = \begin{cases} 1, & \text{prob. } \hat{a} \\ 0, & \text{prob. } 1 - \hat{a}. \end{cases} \tag{2}$$

This last equation shows that, the higher $\hat{a}$ is, the more likely that $\hat{\delta}_t$ will be equal to one. For this reason, we call $\hat{a}$ the insider's *ability*.[7] We assume that nobody (including the insider himself) knows precisely the insider's ability $\hat{a}$ at the outset (i.e. at time zero). Instead, we assume that, a priori, the insider's ability is drawn from the following distribution:

$$\hat{a} = \begin{cases} H, & \text{prob. } \phi_0 \\ L, & \text{prob. } 1 - \phi_0, \end{cases} \tag{3}$$

[6]Throughout the whole paper, we use a "hat" over a variable to denote the fact that it is a random variable.
[7]Equivalently, we could call $\hat{a}$ the insider's information precision.

where $0 < L < H < 1$, and $0 < \phi_0 < 1$. Of course, since the security dividend $\hat{v}_t$ is announced at the end of every period t, the insider will know at the end of every period whether his information for that period was real ($\hat{\delta}_t = 1$) or was just pure noise ($\hat{\delta}_t = 0$).[8] For tractability reasons, we also assume that the market maker observes $\hat{\theta}_t$ at the end of period t, so that his information at the end of every period is the same as the insider's.[9] This information will be useful to both the insider and the market maker in assessing the insider's ability $\hat{a}$.

In making this last assumption, we are essentially saying that the insider's informational advantage over the market maker is one of market timing. Indeed, the insider's information set at the beginning of every period is always exactly one period ahead of the market maker's. Since our goal is not to explain the differences between these two information sets, we reduce the information gap between the insider and the market maker to only (and exactly) one period. Our analysis is then simplified in that both the insider and the market maker perform the same one-period updating at the end of every period, except that the insider's updating will be biased. Preventing the market maker from observing $\hat{\theta}_t$ at the end of period t would simply result in a more complex (non Markov) updating process for the market maker, but would not affect the insider's updating process, in which we are ultimately interested. It would, however, increase the insider's expected profits since the market maker's informational disadvantage would then be greater.

As mentioned above, our model seeks to describe the behavior of an informed trader with a learning bias. In particular we want to model the phenomenon that traders usually think "too much of their ability" when they have been successful at predicting the market in the past. In statistical terms, this will mean that traders update their ability beliefs too much when they are right. Before we formally include this behavior into our model, let us describe how a rational/unbiased insider would react to the information he gathers from past trading rounds.

Let $\hat{s}_t$ denote the number of times that the insider's information was real in the first t periods, that is

$$\hat{s}_t = \sum_{u=1}^{t} \hat{\delta}_u. \tag{4}$$

It can be shown, using Bayes' rule, that, at the end of t periods, a rational insider's updated beliefs

[8] This will be the case since $\hat{\varepsilon}_t = \hat{v}_t$ happens with zero probability with the continuous distributions that we will specify later.

[9] As will become clear below, we could have equivalently assumed that every trader's order and identity (insider vs liquidity trader) are revealed at the end of every period.

about his own ability will be given by

$$\phi_t(s) \equiv \Pr\{\hat{a} = H \mid \hat{s}_t = s\} = \frac{H^s(1-H)^{t-s}\phi_0}{H^s(1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)}. \tag{5}$$

We denote this rational insider's updated expected ability by

$$\mu_t(s) \equiv \mathrm{E}[\hat{a} \mid \hat{s}_t = s] = H\phi_t(s) + L[1 - \phi_t(s)]. \tag{6}$$

In fact, since we do not assume any kind of irrational behavior on the part of the market maker, and since the market maker's information set is the same as the insider's at the *end* of every period, this will be the market maker's updated belief at the end of period t.

In modeling the self-serving attribution bias (which we simply refer to as the *learning bias* from now on), we assume that a trader who successfully forecasts a dividend, weights this success too heavily when applying Bayes' rule to assess his own ability. In choosing our updating rule we seek to accurately model the behavior psychologists have observed, to create a simple, tractable model, and to choose an updating rule that departs gradually from rational updating, and can therefore be used to describe traders with different degrees of bias, including unbiased traders. Psychologists find that when people succeed, they are prone to believe that success was due to their personal abilities rather than to chance or outside factors; when they fail, they tend to attribute their failure to chance and outside factors rather than to their lack of ability. They also find that "self-enhancing attributions for success are more common than self-protective attributions for failure" (Fiske and Taylor, 1991; also see Miller and Ross, 1975). This observed behavior can be modeled by assuming that, when a trader applies Bayes' rule to update his belief about his ability, he overweights his successes, he underweights his failures, and he overweights successes more than he underweights failures. The model is simpler and the qualitative results unchanged if one simply assumes, as we do, that successes are weighted too heavily and failures are weighted correctly.

More precisely, we assume that when evaluating his own ability the insider overweights his successes at predicting the security's dividend (i.e. every time that $\hat{\delta}_t = 1$) by a learning bias factor $\gamma \geq 1$ ($\gamma = 1$ representing a rational insider). For example, at the end of the first period, if the insider finds that $\hat{\theta}_1 = \hat{v}_1$ (i.e. $\hat{\delta}_1 = 1$), the insider will adjust his beliefs to

$$\bar{\phi}_1(1) \equiv \Pr_b\{\hat{a} = H \mid \hat{s}_1 = 1\} = \frac{\gamma H\phi_0}{\gamma H\phi_0 + L(1-\phi_0)}, \tag{7}$$

where the subscript to "Pr" denotes the fact that the probability is calculated by a biased insider. This updated probability is larger than that of a rational insider, i.e. $\bar{\phi}_1(1) \geq \phi_1(1)$. Also, as can be seen from (7), $\bar{\phi}_1(1)$ will be higher the larger γ is, and $\bar{\phi}_1(1) \to 1$ as $\gamma \to \infty$; in other words, the learning bias dictates by how much the insider adjusts his beliefs towards being a high ability insider. Moreover, our model departs continuously from rationality in the sense that $\bar{\phi}_1(1) \to \phi_1(1)$ as $\gamma \to 1$. It is easily shown that, in this case,

$$\bar{\phi}_t(s) \equiv \mathrm{Pr}_b\{\hat{a} = H \mid \hat{s}_t = s\} = \frac{(\gamma H)^s (1-H)^{t-s}\phi_0}{(\gamma H)^s (1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)}, \tag{8}$$

and the (biased) insider's updated expected ability is given by

$$\bar{\mu}_t(s) \equiv \mathrm{E}_b[\hat{a} \mid \hat{s}_t = s] = H\bar{\phi}_t(s) + L[1 - \bar{\phi}_t(s)]. \tag{9}$$

At the beginning of every period t, the risk-neutral insider observes his signal $\hat{\theta}_t$; he then chooses his demand for the risky security in order to maximize his expected period t profits,[10] $\tilde{\pi}_t$, conditional on both his signal and his ability beliefs $\bar{\mu}_{t-1}(\hat{s}_{t-1})$ (which only depend on $\hat{s}_{t-1}$). We denote this demand by

$$\hat{x}_t = X_t(\hat{\theta}_t, \hat{s}_{t-1}). \tag{10}$$

The other trader in the economy is a trader who trades for liquidity purposes in every period. This *liquidity trader*'s demand in period t is given exogenously by the random variable $\hat{z}_t$. Both orders, $\hat{x}_t$ and $\hat{z}_t$, are sent to a market maker who fills the orders. As in Kyle (1985), we assume that the market maker is risk-neutral and competitive, and will therefore set prices so as to make zero expected profits. So, if we denote the total order flow coming to the market maker in period t by

$$\hat{\omega}_t = \hat{x}_t + \hat{z}_t, \tag{11}$$

[10] As we mention above, both the risky dividend and the insider's signal are announced at the end of every period, so that the market maker is always exactly one period behind the insider in terms of information at the beginning of the next period. This implies that the insider never finds it optimal to suboptimally choose his demand for one period in order to maximize longer-term profits.

the market maker will set the security's price equal to

$$\hat{p}_t = P_t(\hat{\omega}_t, \hat{s}_{t-1}) \equiv \mathrm{E}[\hat{v}_t \mid \hat{\omega}_t, \hat{s}_{t-1}] \tag{12}$$

in period t. An equilibrium to our model is defined as a sequence of pairs of functions (X_t, P_t), $t = 1, 2, \ldots,$ such that the insider's demand in period t, $X_t(\hat{\theta}_t, \hat{s}_{t-1})$ maximizes his expected profits (according to his own beliefs) for that period given that he faces a price curve P_t, while the market maker is expecting zero profit in that period.

As it will become obvious later, the main results of this paper are driven by the insider's updating dynamics. As such, the liquidity trader and the market maker only play minimal roles in this model. In fact, their role is essentially one of market clearing. Indeed, the presence of the liquidity trader introduces noise that will prevent the "no trade equilibrium" described by Milgrom and Stokey (1982) from occurring. The competitive market maker assumption is simply made out of convenience; the presence of a risk-neutral rational trader would serve a similar purpose, as in Daniel, Hirshleifer and Subrahmanyam (1998). Appendix D presents a brief synopsis of alternative specifications that in turn avoid the presence of the liquidity trader and the market maker, introduce the possibility for the market maker to maximize profits, and increase the number of market participants. As argued in that appendix, these specifications do not affect the nature of the insider's (or insiders') updating and, therefore, do not change the main results of the paper.

2 A Linear Equilibrium

In this section, we show that, when $\hat{v}_t$, $\hat{\varepsilon}_t$, and $\hat{z}_t$ are jointly and independently normal, there is a linear equilibrium to our economy. We use that linear equilibrium in section 3 to illustrate the properties of the model. More precisely, for this and the next section, we assume that

$$\begin{bmatrix} \hat{v}_t \\ \hat{\varepsilon}_t \\ \hat{z}_t \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma & 0 & 0 \\ 0 & \Sigma & 0 \\ 0 & 0 & \Omega \end{bmatrix} \right), \ t = 1, 2, \ldots, \tag{13}$$

and that each such vector is independent of all the others. Note that it is crucial that $\mathrm{Var}(\hat{v}_t) = \mathrm{Var}(\hat{\varepsilon}_t)$, since we do not want the size of $\hat{\theta}_t$ to reveal anything about the likelihood that $\hat{\delta}_t = 1$ until $\hat{v}_t$ is announced. In other words, ability updating is only possible when both $\hat{\theta}_t$ and $\hat{v}_t$ become observable.

Let us conjecture that, in equilibrium, the function $X_t(\theta, s)$ is linear in θ, and that the function $P_t(\omega, s)$ is linear in ω:

$$\begin{aligned} X_t(\theta, s) &= \beta_t(s)\,\theta, &\text{(14a)}\\ P_t(\omega, s) &= \lambda_t(s)\,\omega. &\text{(14b)}\end{aligned}$$

Our objective is to find $\beta_t(s)$ and $\lambda_t(s)$ which are consistent with this conjecture. We start with the following result.

Lemma 2.1 *Assume that a linear equilibrium exists in period t, that is assume that (14a) and (14b) hold. Then, in period t, the insider's demand for the risky security is given by*

$$\hat{x}_t = \frac{\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t}{2\lambda_t(\hat{s}_{t-1})}, \tag{15}$$

and the market maker's price schedule is given by

$$\hat{p}_t = \frac{\mu_{t-1}(\hat{s}_{t-1})\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t. \tag{16}$$

Proof: See Appendix A.

This lemma establishes that we can indeed write $\hat{x}_t = \beta_t(\hat{s}_{t-1})\hat{\theta}_t$ with

$$\beta_t(s) = \frac{\bar{\mu}_{t-1}(s)}{2\lambda_t(s)}, \tag{17}$$

and $\hat{p}_t = \lambda_t(\hat{s}_{t-1})\hat{\omega}_t$ with

$$\lambda_t(s) = \frac{\mu_{t-1}(s)\beta_t(s)\Sigma}{\beta_t^2(s)\Sigma + \Omega}. \tag{18}$$

However, the result relies on the assumption that a linear equilibrium exists. It turns out that this assumption is not always satisfied given the insider's learning bias. In fact, the following lemma derives the exact condition under which such an equilibrium will exist in a given period t.

Lemma 2.2 *In any given period t, there exists a linear equilibrium of the form conjectured in (14a) and (14b) if and only if $\bar{\mu}_{t-1}(\hat{s}_{t-1}) \leq 2\mu_{t-1}(\hat{s}_{t-1})$.*

Proof : See Appendix A.

This condition essentially states that the biased insider's beliefs about his ability cannot exceed those of the rational market maker by too much for an equilibrium to exist. It effectively ensures that the insider will on average be making profits, even though he may not be optimizing correctly due to his biased beliefs.[11] Were this condition not satisfied, the market maker would know that the insider will take a position that will on average result in negative profits. Our condition that the competitive market maker quotes a price schedule earning him on average zero profits can then never be satisfied. This results in a market breakdown. The following lemma derives a condition that is both necessary and sufficient to always avoid such outcomes.

Lemma 2.3 *A necessary and sufficient condition for $\bar{\mu}_{t-1}(\hat{s}_{t-1}) > 2\mu_{t-1}(\hat{s}_{t-1})$ to be avoided for any history up to any period t and any $\gamma > 1$ is that $H \leq 2L$.*

Proof : See Appendix A.

Note that, for a given *fixed* value of γ, $H \leq 2L$ is too strong a condition to avoid market breakdowns for any history, i.e. the condition is then sufficient but not necessary. However, since market breakdowns are outside the scope of this paper, we assume that $H \leq 2L$ is always satisfied in the rest of our analysis.[12] Such an assumption allows us to vary γ throughout the paper without having to worry about the existence of an equilibrium, but does not affect the qualitative aspects of our results.[13] We finish this section with the following characterization of the equilibrium.

Lemma 2.4 *Assuming that $H \leq 2L$, there is always a unique linear equilibrium to the economy described in section 1. In this equilibrium, the insider's demand and the market maker's price*

[11] Note that this condition does not violate the absence of arbitrage results established by Bossaerts (1999) for security prices, and by De Finetti (1937) for subjective probabilities. Indeed, although the insider's learning bias could introduce profit opportunities for other informed traders (were they present in the economy), they do not lead to riskless arbitrage opportunities. In fact, the beliefs of the market maker in our model correspond to the *market beliefs* in Bossaerts' work. Note also that, if the insider is risk-averse, this existence condition is not needed; however, the model then loses tractability.

[12] For a formal analysis of market breakdowns, consult Bhattacharya and Spiegel (1991).

[13] In fact, although this assumption guarantees the survival of the insider by making his expected profits positive, it does not guarantee the survival of overconfidence for an individual, as we show in section 3.2.

schedule are given by (14a) and (14b) with

$$\beta_t(s) = \sqrt{\frac{\Omega}{\Sigma}\frac{\bar{\mu}_{t-1}(s)}{2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)}}, \quad \textit{and} \tag{19a}$$

$$\lambda_t(s) = \frac{1}{2}\sqrt{\frac{\Sigma}{\Omega}\bar{\mu}_{t-1}(s)\left[2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)\right]}. \tag{19b}$$

Proof : See Appendix A.

In what follows, we use this equilibrium to study the effects of the insider's learning bias on the economy.

3 Properties of the Model

In this section, we analyze the effects of the insider's learning bias on the properties and dynamics of the economy in equilibrium. We introduce a measure of overconfidence analogous to that found in static models and show that the learning bias results in dynamically evolving insider overconfidence. We then look at the effect of this changing overconfidence on trading volume, trader profits, price volatility, as well as the expected price patterns.

3.1 Convergence

If this financial market game is played to infinity, we would expect both the insider and the market maker to eventually learn the exact ability $\hat{a}$ of the insider. This in fact would be true for a rational insider ($\gamma = 1$). However, since our insider learns his ability with a personal bias, this result is not immediate; in fact, as we shall see, this result is not true for a highly biased insider.

When $\hat{a} = H$, we expect the insider to correctly guess the one-period dividend a fraction H of the time. So, as we play the game more and more often (as t tends to ∞), we expect his updated posteriors

$$\bar{\phi}_t(s) = \frac{(\gamma H)^s(1-H)^{t-s}\phi_0}{(\gamma H)^s(1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)} = \frac{1}{1+\left(\frac{L}{\gamma H}\right)^s\left(\frac{1-L}{1-H}\right)^{t-s}\frac{1-\phi_0}{\phi_0}}$$

to behave like

$$\frac{1}{1+\left(\frac{L}{\gamma H}\right)^{Ht}\left(\frac{1-L}{1-H}\right)^{t-Ht}\frac{1-\phi_0}{\phi_0}} = \frac{1}{1+\left[\left(\frac{L}{\gamma H}\right)^{H}\left(\frac{1-L}{1-H}\right)^{1-H}\right]^{t}\frac{1-\phi_0}{\phi_0}}.$$

This last quantity will converge to 1 as desired if

$$\left[\left(\frac{L}{\gamma H}\right)^{H}\left(\frac{1-L}{1-H}\right)^{1-H}\right]^{t} \to 0 \quad \text{as } t \to \infty,$$

or equivalently if

$$\left(\frac{L}{\gamma H}\right)^{H}\left(\frac{1-L}{1-H}\right)^{1-H} < 1.$$

The following lemma shows that this is indeed the case.

Lemma 3.1 *When $\hat{a} = H$, the updated posteriors of the insider $\bar{\phi}_t(\hat{s}_t)$ will converge to 1 almost surely as $t \to \infty$.*

Proof : See Appendix B.

So both the insider and the market maker will eventually learn the insider's ability precisely when it is high (when $\hat{a} = H$). Let us now turn to the case where $\hat{a} = L$. In this case, we expect the insider to correctly guess the one-period dividend a fraction L of the time. So, as we play the game more and more often (as t tends to ∞), we expect his updated posteriors $\bar{\phi}_t(s)$ to behave like

$$\frac{1}{1+\left(\frac{L}{\gamma H}\right)^{Lt}\left(\frac{1-L}{1-H}\right)^{t-Lt}\frac{1-\phi_0}{\phi_0}} = \frac{1}{1+\left[\left(\frac{L}{\gamma H}\right)^{L}\left(\frac{1-L}{1-H}\right)^{1-L}\right]^{t}\frac{1-\phi_0}{\phi_0}}.$$

As the following lemma shows, this quantity only converges to zero when γ is close enough to 1. This means that the market maker will always learn the insider's ability when it is low (when $\hat{a} = L$),[14] but the insider will only do so if his learning bias is not too large.

[14]This is due to the fact that we assume that the market maker's learning is unbiased.

Lemma 3.2 *When $\hat{a} = L$, the updated posteriors of the insider $\bar{\phi}_t(\hat{s}_t)$ will converge as follows:*

$$\bar{\phi}_t(\hat{s}_t) \xrightarrow{a.s.} \begin{cases} 0, & \textit{if } \gamma < \gamma^* \\ \phi_0, & \textit{if } \gamma = \gamma^* \\ 1, & \textit{if } \gamma > \gamma^*, \end{cases}$$

where $\gamma^* = \frac{L}{H}\left(\frac{1-L}{1-H}\right)^{(1-L)/L}$.

Proof : See Appendix B.

One implication of this lemma is that a low ability insider whose learning bias is sufficiently extreme may never acknowledge his low ability, no matter how much experience he has.[15] To illustrate this, we show in Figures 1(a) and 1(b) how we expect an insider to adjust his beliefs about his own ability ($\bar{\phi}_t$) when his actual ability is high, and when it is low respectively. As seen in these figures, the biased insider's beliefs are always on average larger than those of an unbiased but otherwise identical insider. Since unbiased insiders always eventually learn their ability, it is therefore not surprising to find that high ability insiders also always learn their own ability:[16] they naturally tend to update towards that high ability. However, as shown in Figure 1(b), a biased insider may not always give in to his observations: more precisely, if $\gamma > \gamma^*$, he will never find out if he is a low ability insider.

3.2 Patterns of Overconfidence

As we show in section 3.1, the insider will eventually learn his own ability, provided that his learning is not too biased (i.e. provided that γ is not too large). This means that, when the insider's ability is low, the insider eventually comes to his senses, and recognizes the fact that he is a low ability insider. However, it is always the case that the insider thinks too highly of himself relative to an otherwise identical unbiased insider. This section introduces a measure for this discrepancy; we call it the insider's *overconfidence*. The evolution of the insider's overconfidence throughout his life is central to our study, as this overconfidence essentially measures by how much our model departs from a purely rational setup in any given period.

[15] Bossaerts (1999) shows that prior beliefs that are correctly updated using Bayes' law converge to the right posterior beliefs, whether the priors are biased or not. In contrast, our result shows that correct priors updated incorrectly may not lead to the correct posteriors.

[16] In fact, they will do so faster the more biased they are.

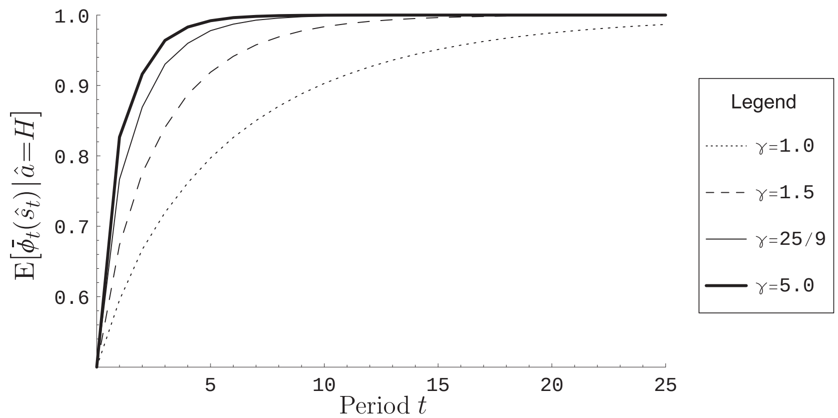

(a) Convergence when $\hat{a} = H$.

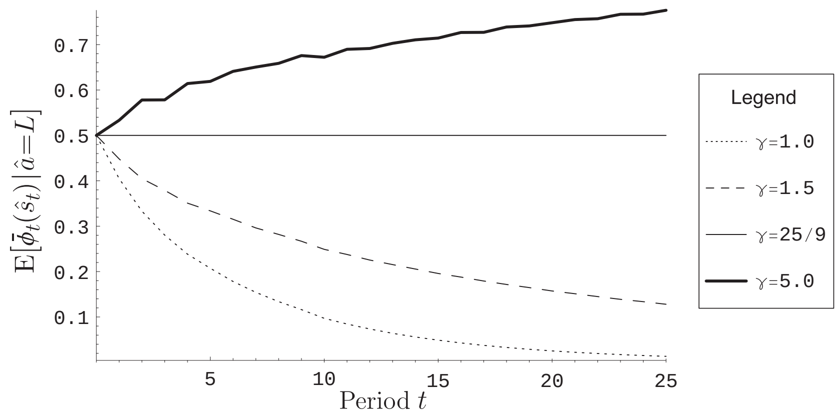

(b) Convergence when $\hat{a} = L$.

Figure 1: Convergence patterns of the insider's expected beliefs about his own ability when (a) $\hat{a} = H$; (b) $\hat{a} = L$. Both figures were obtained with $H = 0.9$, $L = 0.5$, $\phi_0 = 0.5$, and $\Sigma = \Omega = 1$. Each line was drawn with a different γ, shown in the legends. Note that, with these parameter values, $\gamma^* = 25/9 \approx 2.78$.

In our model, an insider is considered very overconfident at the end of a particular period t if his updated expected ability at that time ($\bar{\mu}_t(\hat{s}_t)$) is large compared to the updated expected ability that a rational insider would have reached with the same past history of successes and failures ($\mu_t(\hat{s}_t)$). To measure the insider's overconfidence at the end of t periods, we therefore define the random variable

$$\hat{\kappa}_t \equiv K_t(\hat{s}_t) \equiv \frac{\bar{\mu}_t(\hat{s}_t)}{\mu_t(\hat{s}_t)}. \tag{20}$$

Of course, when the insider is rational ($\gamma = 1$), the numerator is exactly equal to the denominator of this expression, and $\hat{\kappa}_t = 1$ for all $t = 1, 2, \ldots$. On the other hand, when the insider's learning is biased ($\gamma > 1$), we have $\bar{\mu}_t(\hat{s}_t) \geq \mu_t(\hat{s}_t)$, and $\hat{\kappa}_t \geq 1$ for all t. As the next proposition shows, the insider's overconfidence in period t is greater when the insider's learning bias is large. In other words, the insider's overconfidence is directly attributable to his learning bias.

Proposition 3.1 *The function $K_t(s)$ defined in (20) is increasing in γ.*

Proof: See Appendix B.

Our measure of overconfidence at any point in time is therefore increasing in the insider's learning bias, but is it also increasing in the number of his past successful predictions? Since the insider's overconfidence results from his learning bias when he is successful, it is tempting to conclude that the more successful an insider is, the more overconfident he will be. As we next show, this intuition is wrong.

First, since the insider updates his beliefs incorrectly only after successful predictions, it is always true that $\bar{\mu}_t(0) = \mu_t(0)$, and therefore $K_t(0) = 1$. However, as soon as the insider successfully predicts one risky dividend, his learning bias makes him overconfident,[17] and $\bar{\mu}_t(1) > \mu_t(1)$, so that $K_t(1) > 1$. So, it is always true that the insider's first successful prediction makes him overconfident.[18] However, it is not always the case that an additional successful prediction always makes the insider more overconfident.

To see this, suppose that we are at the end of the second period. The insider will then have been successful 0, 1 or 2 times. We already know that $K_2(1) > K_2(0) = 1$ for any value of the insider's learning bias parameter γ. Now, suppose that γ is large. This means that if the insider is successful in the first period, he will immediately (and perhaps falsely) jump to the conclusion that he is a high ability insider, i.e. $\bar{\mu}_1(1)$ is close to H. Since this one successful period has already convinced the insider that his ability is high, the second period results will not have much of an effect on his beliefs, whether he is successful or not in that period, i.e. $\bar{\mu}_2(2)$ is close to $\bar{\mu}_2(1)$. On the other hand, if the insider had been rational ($\gamma = 1$), he would have adjusted his expected ability beliefs more gradually. In particular, after a first period success, a rational insider does

[17] In our model, traders are not overconfident when they begin to trade. It is through making forecasts and trading that they become overconfident. This leads market participants to be, on average, overconfident. In real markets, selection bias may cause even beginning traders to be overconfident. Indeed, since not everyone trades, it is likely that people who rate their own trading abilities most highly will seek jobs as traders or will trade actively on their own account. Those with actual high ability and those with high overconfidence will rate their own ability highest. Thus, even at the entry level, we would expect to find overconfident traders.

[18] Also, as section 3.1 shows, he will remain so for the rest of his life.

not automatically conclude that his ability is high. Instead, he adjusts his posterior expected ability beliefs towards H, and uses the second period result to further adjust these beliefs: upward towards H if he is successful, and downward towards L otherwise. As a result, $\mu_2(2)$ will be somewhat larger than $\mu_2(1)$. Therefore, since $\bar{\mu}_2(2) \approx \bar{\mu}_2(1)$ and $\mu_2(2) > \mu_2(1)$, we have

$$K_2(2) \equiv \frac{\bar{\mu}_2(2)}{\mu_2(2)} < \frac{\bar{\mu}_2(1)}{\mu_2(1)} \equiv K_2(1),$$

and we see that $K_2(s)$ decreases when s goes from 1 to 2.

In short, the biased insider adjusts his beliefs non-rationally with every successful prediction, making him overconfident. However, when the insider's past performance is sufficiently good (large number of successful predictions), it is the case that even an unbiased insider would reach the conclusion that he is a high ability insider. In other words, the significance of the insider's past performance overweights the significance of his learning bias. The following proposition describes this phenomenon in more details.

Proposition 3.2 *The function $K_t(s)$ defined in (20) is increasing over $s \in \{0, \ldots, s_t^*\}$ and decreasing over $s \in \{s_t^*, \ldots, t\}$, for some $s_t^* \in \{1, \ldots, t\}$.*

Proof : See Appendix B.

Intuitively, this result says that a trader who has been very successful in only a few rounds of trading or one who has been moderately successful in several rounds of trading will have a greatly inflated opinion of his ability. But a trader who has been very successful over many rounds of trading probably does have high ability. And while he may overestimate his expected ability, he does not do so by as much as do moderately successful traders.

In this model, traders are rational in all respects except that they have a common learning bias: they tend to attribute their successes disproportionately to their own ability. This leads successful traders to become overconfident. Other learning biases can also lead to overconfidence. For example, it is well known that, when updating beliefs from sequential information, people tend to weight recent information too heavily and older information too little (Anderson, 1959, 1981; Hogarth and Einhorn, 1992).[19] We do not introduce this recency effect into our model because

[19] In experimental studies, subjects sometimes also exhibit a "primacy effect," weighting the earliest observations of a time series heavily (Anderson, 1981). This happens most often in situations where subjects lose interest in the data (Hogarth and Einhorn, 1992). It is unlikely that traders would lose interest in their own successes and failures, and so we would not expect to find a large primacy effect in their updating.

doing so would negate the Markov property of the insider's (and the market maker's) updating process. It is clear though that, if traders weight recent outcomes more heavily than older ones, recently successful traders will tend to become overconfident.

The last two propositions describe how the insider's overconfidence in a particular period depends on his learning bias and on his previous performance. Let us now turn to how his overconfidence is expected to behave over time. To do this, we calculate the ex ante expected period t overconfidence level of the insider, $\mathrm{E}[\hat{\kappa}_t]$. Since

$$\begin{aligned}\Pr\{\hat{s}_t = s\} &= \Pr\{\hat{s}_t = s \mid \hat{a} = H\} \Pr\{\hat{a} = H\} + \Pr\{\hat{s}_t = s \mid \hat{a} = L\} \Pr\{\hat{a} = L\} \\ &= \binom{t}{s} H^s (1-H)^{t-s} \phi_0 + \binom{t}{s} L^s (1-L)^{t-s} (1-\phi_0), \end{aligned} \tag{21}$$

we have

$$\mathrm{E}[\hat{\kappa}_t] = \sum_{s=0}^{t} \binom{t}{s} \left[H^s (1-H)^{t-s} \phi_0 + L^s (1-L)^{t-s} (1-\phi_0) \right] \frac{\bar{\mu}_t(s)}{\mu_t(s)}, \tag{22}$$

where $\mu_t(s)$ and $\bar{\mu}_t(s)$ are as in (6) and (9). Figure 2 shows the patterns in the expected level of overconfidence for different values of γ. When γ is relatively small ($\gamma < \gamma^*$), the insider will on average be overconfident at first but, over time, will converge to a rational behavior.

This can be explained as follows. Over a small number of trading periods a trader's success rate may greatly exceed that predicted by his ability. Very successful traders will overestimate the likelihood that success is due to ability rather than luck. But over many trading periods a trader's success rate is likely to be close to that predicted by his ability. Only those traders with extreme learning bias (or with very unlikely success patterns) will fail to recognize their true ability. Indeed, as γ increases, the insider tends to put more and more weight on his past successes, and so takes a little more time to rationally find his ability. However, if γ is too large (more precisely, if $\gamma > \gamma^*$), it is possible that the insider puts so much weight on his past successes in the stock market that he never becomes rational. It can be shown that $\mathrm{E}[\hat{\kappa}_t]$ then converges to $\phi_0 + (1-\phi_0)\frac{H}{L}$.

Thus our model predicts that more inexperienced traders will be more overconfident than experienced traders. Less experienced traders are more likely to have success records which are unrepresentative of their abilities. For some, this will lead to overconfidence. By the law of large numbers, older traders are likely to have success records which are more representative of their abilities; they will, on average, have more realistic self assessments. In a sufficiently large group

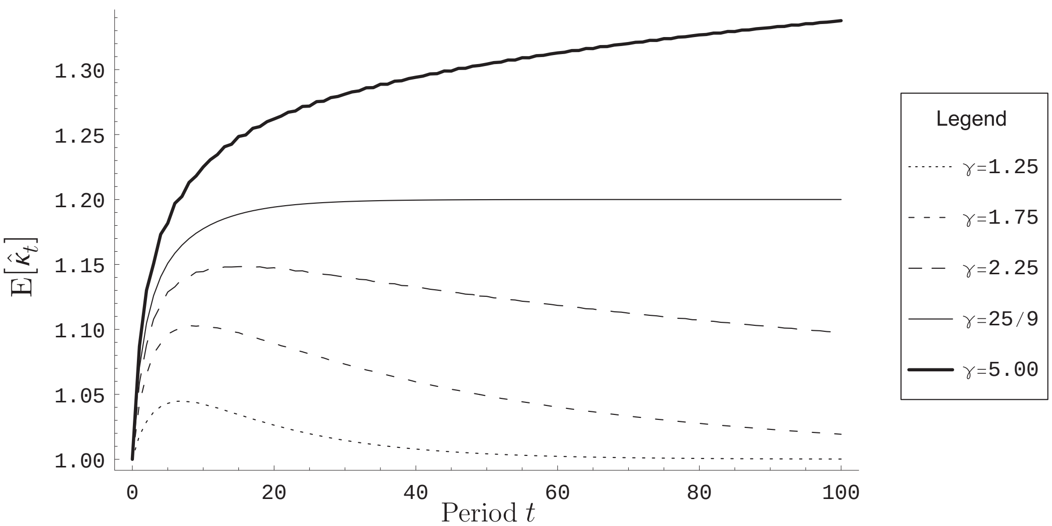

Figure 2: Ex ante expected patterns in the level of overconfidence of the insider over time. The figure was obtained with $H = 0.9$, $L = 0.5$, $\phi_0 = 0.5$, $\Sigma = \Omega = 1$. Each line was obtained with a different value of γ shown in the legend. Note that, with these parameter values, $\gamma^* = 25/9 \approx 2.78$.

of traders,[20] however, there will be some successful, older, low-ability traders with whom the odds have not yet caught up. These traders are likely to make large mistakes in the future.

As we see in Figure 2, on average, a trader's overconfidence increases during the first part of his trading experience and decreases during the latter part. It is natural to ask what factors determine the point at which a trader's overconfidence is likely to peak. All other things being equal, the greater a trader's learning bias, γ, the longer it is likely to take for his overconfidence level to peak. Figure 3 illustrates this effect.

In addition to the degree of learning bias, how quickly a trader's overconfidence peaks (and how quickly he ultimately learns his true ability) depends on the frequency, speed, and clarity of the feedback he receives. A trader who receives frequent, immediate, and clear feedback will, on average, peak in overconfidence early, and quickly realize his true ability. One who receives infrequent, delayed, and ambiguous feedback will peak in overconfidence later, and more slowly realize his true ability. In general, financial markets are difficult environments for learning, as feedback is often ambiguous and comes well after a decision was made. We would expect those who trade most frequently, such as professional traders, and those who keep careful records rather than relying on memory to learn most quickly. In our model a trader receives immediate feedback each period. The greater the difference between the high (H) and the low (L) ability levels, the clearer

[20] Appendix D.3 presents an extension of the model with multiple insiders. The learning and overconfidence dynamics are essentially the same as in this single insider model.

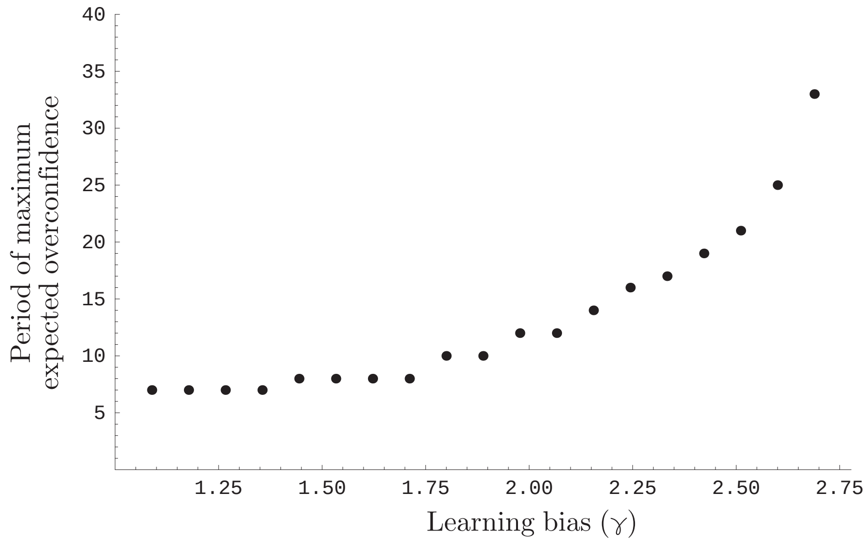

Figure 3: Period of maximum expected overconfidence as a function of the insider's learning bias γ. The figure was obtained with $H = 0.9$, $L = 0.5$, and $\phi_0 = 0.5$.

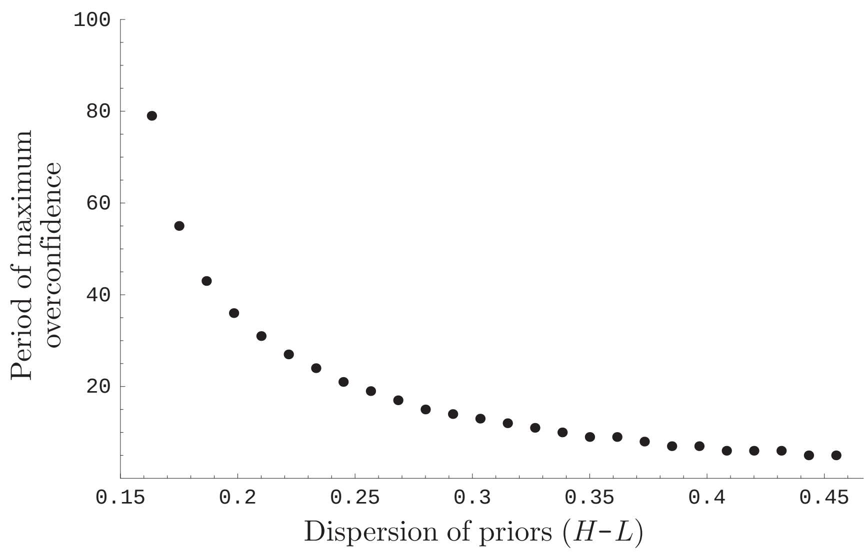

Figure 4: Period of maximum expected overconfidence as a function of the dispersion $H - L$ of the insider's prior ability beliefs. The figure was obtained with $\phi_0 = 0.5$, $\gamma = 1.1$, and keeping the insider's ex ante expected ability $\phi_0 H + (1 - \phi_0)L$ constant at 0.7.

the feedback. Figure 4 illustrates how long it takes on average for the insider's overconfidence to peak as a function of $H - L$.

3.3 Effects of the Insider's Learning Bias

Section 3.2 shows that the insider's learning bias has a dynamic impact on his beliefs about his ability and on the way he interprets future private information. This in turn affects the future trading process. In this and the next section, we describe how this trading process, as measured by trading volume, trader profits, and price volatility, is affected.

Since the insider's learning bias is unaffected by his success rate and vice versa,[21] our model allows us to analyze the effects of the learning bias on the insider's behavior and on the properties of the economy in two different ways. First, given a fixed past history of the insider's successes and failures, we can vary the size of his learning bias to get an idea of the impact of that bias. This is the focus of the current section. Second, we can fix the insider's learning bias, and determine the effects of different trading histories on the insider and the economy in general. We will turn to this in section 3.4.

Before embarking on the effects of the insider's learning bias, we calculate the three quantities that will help us measure these effects: trading volume, expected insider profits, and price volatility.[22] Let $\hat{\psi}_t$ denote the trading volume in period t. Since this trading volume comes from both the insider and the liquidity trader, it is given by

$$\hat{\psi}_t \equiv \frac{1}{2}\left(|\hat{x}_t| + |\hat{z}_t|\right). \tag{23}$$

The following lemma shows how the expected one-period trading volume, expected insider profits and price volatility are calculated, conditional on the insider having been successful s times in the first t periods.

Lemma 3.3 *Conditional on the insider having been successful s times in the first t periods (i.e. conditional on $\hat{s}_t = s$), the expected volume, the expected insider profits, and the price variance in period $t+1$ are given by*

$$\mathrm{E}\left[\hat{\psi}_{t+1} \mid \hat{s}_t = s\right] = \frac{1}{\sqrt{2\pi}}\left[\sqrt{\Sigma}\,\beta_{t+1}(s) + \sqrt{\Omega}\right], \tag{24}$$

[21]The learning bias γ is constant in every period, i.e. it does not change based on the insider's past successes. Also, the insider's successes, on which his updating is correctly based, are the direct results of his ability whatever the insider's learning bias may be, i.e. his success rate (unlike his profits) is not affected by his learning bias.

[22]We will study the effects of the learning bias on the patterns of prices in a separate section 3.5.

$$\mathrm{E}\left[\hat{\pi}_{t+1} \mid \hat{s}_t = s\right] = \frac{1}{2}\sqrt{\Sigma\Omega}\,\sqrt{\bar{\mu}_t(s)\left[2\mu_t(s) - \bar{\mu}_t(s)\right]}, \tag{25}$$

and

$$\mathrm{Var}\left[\hat{p}_{t+1} \mid \hat{s}_t = s\right] = \frac{\Sigma}{2}\,\bar{\mu}_t(s)\,\mu_t(s), \tag{26}$$

respectively.

Proof : See Appendix B.

Recall from equation (14a) in section 2 that the insider will multiply his period $t+1$ signal, $\hat{\theta}_{t+1}$, by $\beta_{t+1}(s)$ to obtain his demand for the risky asset in that period. In other words, $\beta_{t+1}(s)$ represents the insider's trading intensity in period $t+1$ after having been successful s times in the first t periods. As the above lemma shows, greater average insider intensity leads to larger expected volume.

Moreover, as shown in section 3.2, a biased insider, who has had at least one success, is always overconfident. In other words, the insider thinks that his signal $\hat{\theta}_{t+1}$ in period $t+1$ is more informative than it really is. This leads him to use his information more aggressively than he should and leads to higher expected trading volume in the risky security. As the next proposition demonstrates, the greater the learning bias the greater this trading.

Proposition 3.3 *Given that $\hat{s}_t = s$, the expected volume in period $t+1$ is increasing in the insider's learning bias parameter γ.*

Proof : See Appendix B.

As mentioned above, expected volume in a particular period will be larger the larger the expected insider trading intensity is for that period. Notice that we can rewrite $\beta_{t+1}(s)$ given in (19a) as

$$\beta_{t+1}(s) = \sqrt{\frac{\Omega}{\Sigma}\left[\frac{2}{K_t(s)} - 1\right]^{-1}}, \tag{27}$$

where $K_t(s)$ is as defined in equation (20). This tells us that the trading intensity $\beta_{t+1}(s)$ of the insider in period $t+1$ given that he has been successful s times in the first t periods is a monotonically increasing function of the insider's overconfidence $K_t(s)$ after t periods. Since we showed in Proposition 3.1 that the insider's overconfidence in any period, given any number of past successes, is increasing in γ, it is natural to find that expected volume in a particular period will also be increasing in γ.

Let us now look at the effect of the learning bias on the insider's profits. We know that the biased insider trades too aggressively on his information; in other words, the insider's learning bias makes him act suboptimally. It is therefore not surprising that the insider's expected profits in any given period are decreasing in his learning bias parameter γ.

Proposition 3.4 *Given that* $\hat{s}_t = s$*, the expected insider profits in period* $t+1$ *are decreasing in the insider's learning bias parameter* γ*.*

Proof: See Appendix B.

The more overconfident the insider, the more he trades in response to any given signal. This increases his expected trading relative to that of the liquidity trader. Therefore the signal to noise ratio in total order flow increases and the market-maker is able to make better inferences about the insider's signal. The market-maker is then able to set prices that vary more in response to $\hat{\theta}_t$ and are closer to the expected dividend conditional on the insider's signal ($\mathrm{E}[\hat{v}_t \mid \hat{\theta}_t]$) and further from its unconditional expectation (zero). This increases price volatility.

Proposition 3.5 *Given that* $\hat{s}_t = s$*, the expected price volatility in period* $t+1$*, as measured by the price's variance, is increasing in the insider's learning bias parameter* γ*.*

Proof: See Appendix B.

3.4 Effects of the Insider's Past Performance

The effects described in section 3.3 are static in the sense that they do not depend on the insider's past performance. Given any success history, the next period's trading volume and price volatility are expected to be larger and the next period's insider profits are expected to be lower when γ is large. These results are analogous to the results documented by Odean (1998) who shows that trader overconfidence has these effects in a one-period economy. Since, as documented in section 3.2, the insider's learning bias eventually makes him overconfident, our results are natural extensions of Odean's static results.

However, we also know from section 3.2 that the insider's overconfidence changes dynamically with his past performance. In this section we look at how, for an insider with a particular learning bias, the economy is affected by that insider's past performance.

The monotonic relationship between $\beta_{t+1}(s)$ and $K_t(s)$ described in (27) also helps us characterize the conditional expected volume in a particular period $t+1$, given different past histories at the end of t periods. It is not surprising to find that the expected one-period volume given s insider successes in the first t periods has the same shape as $K_t(s)$ as a function of s, which we described in Proposition 3.2.

Proposition 3.6 *The expected volume in period $t+1$, conditional on the insider having been successful s times in the first t periods (i.e. given $\hat{s}_t = s$), is increasing over $s \in \{0, \ldots, s_t^\circ\}$ and decreasing over $s \in \{s_t^\circ, \ldots, t\}$, for some $s_t^\circ \in \{1, \ldots, t\}$.*

Proof : See Appendix B.

In section 3.3, we saw that the insider's profits are reduced by his learning bias. As we shall see next, this learning bias can cause a successful insider's expected future profits to be smaller than a less successful insider's. This is because two forces affect an insider's expected future profits: his overconfidence and his expected ability.

To disentangle these two forces, let us describe the insider's expected profits in period $t+1$ after he has been successful s times in the first t periods. We know from section 3.2 that the insider's overconfidence at the end of t periods is at a minimum of 1 when $s = 0$. We also know from Proposition 3.2 that overconfidence increases with the number s of past insider successful dividend predictions (up to s_t^*). This means that the insider's decision in period $t+1$ will be more and more distorted as s increases.[23]

At the same time, as s increases, it becomes increasingly likely that the insider's ability is high,[24] though not as likely as the insider thinks. A biased insider who becomes sufficiently overconfident may act so suboptimally that he more than offsets the potential increase in expected profits coming from his probably higher ability. As the insider's overconfidence comes back down ($s > \hat{s}_t^*$), successes decrease the insider's overconfidence while increasing his expected ability. Thus both forces lead to additional expected future profits.

Figures 5(a) and 5(b) illustrate how the insider's overconfidence and expected ability counterbalance each other. In Figure 5(a), we look at the insider's expected profits in period 11, as a

[23]This was seen to be true in equation (27), where we show that $\beta_{t+1}(s)$ is monotonically increasing in $K_t(s)$.

[24]Result C.1 of Appendix C shows that more past successes increase the likelihood $\bar{\phi}_t(s)$ that the insider's ability is high.

function of the number of successes he has had in the first 10 periods; we do this for three different values of γ (1, 2, and 5), and otherwise use the same parameters as in Figures 1 and 2. It is clear from that figure that an unbiased insider always benefits from an additional past success, since his expected ability is higher. However, when the insider's learning is biased, it is possible that his overconfidence (which we plot in Figure 5(b)) prevents him from benefiting from the boost in expected ability that results from an additional success. In fact, for this example, we can see that an insider with $\gamma = 2$ or $\gamma = 5$ who has had six successes in the first 10 periods does worse than an insider who has yet to predict one dividend correctly! This simple numerical example can in fact be generalized as follows.

Proposition 3.7 *Given that $\hat{s}_t = s$, the expected insider profits in period $t+1$ are increasing over $s \in \{0, \ldots, s'_t\}$ and $s \in \{s''_t, \ldots, t\}$, but are decreasing over $s \in \{s'_t, \ldots, s''_t\}$ for some $(s'_t, s''_t) \in \{1, \ldots, t\}^2$ such that $s'_t \leq s''_t$.*

Proof: See Appendix B.

Since the insider can only be perfectly right or completely wrong in any given period, the correct measure of his past performance at the beginning of period $t+1$ in this model is the number of his past successes ($\hat{s}_t$). In reality, traders can be right or wrong to different extents, and so the measure that is typically used to measure their performance is past profits. It is easily shown that, in our model, expected past profits are monotonically increasing in the number of past successes. Figure 5(c) illustrates this for the above numerical example. Therefore future expected insider profits, as a function of past profits, are first increasing, then decreasing, and then increasing again.

Corollary 3.1 *Conditional on the insider's cumulative profits π in the first t periods, the expected insider profits in period $t+1$ are increasing over $\pi \in (-\infty, \pi'_t]$ and $\pi \in [\pi''_t, \infty)$, but are decreasing over $\pi \in [\pi'_t, \pi''_t]$ for some $(\pi'_t, \pi''_t) \in \mathbb{R}^2$ such that $\pi'_t \leq \pi''_t$.*

In our model traders trade on their own account. We do not model the agency issues associated with money managers investing for others nor do we model the relationship between an individual money manager and the fund for which he may work. These issues may greatly influence money managers' decisions. Nevertheless, Proposition 3.7 along with Figure 5 may provide some guidance about the choice of managers. We show that a trader with more past successes may have lower expected future profits than a trader with fewer successes. This is because the more successful trader, though objectively more likely to possess high ability, will not make maximum use of his ability due to his overconfidence. An investor choosing a money manager cannot usually observe

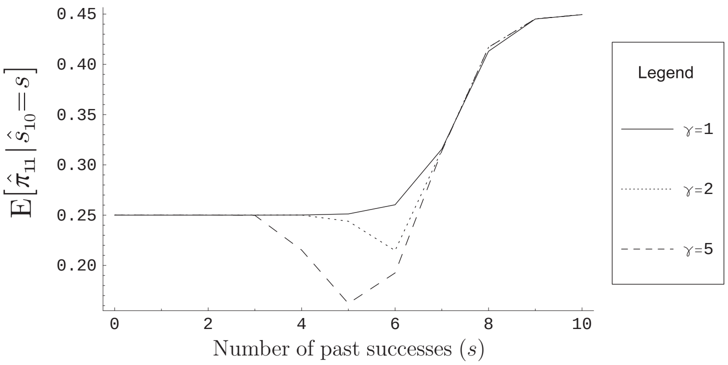

(a) Expected insider profits.

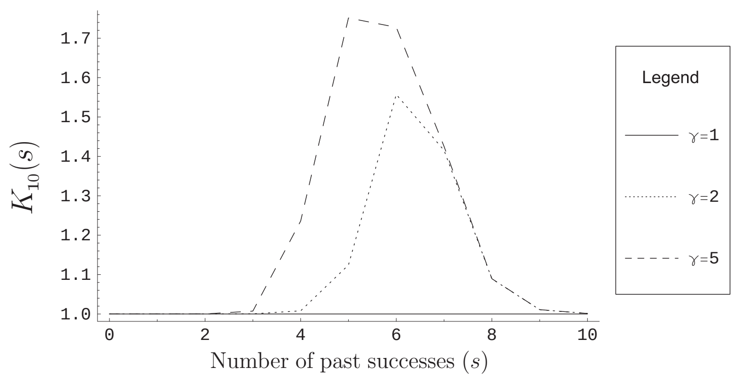

(b) Overconfidence.

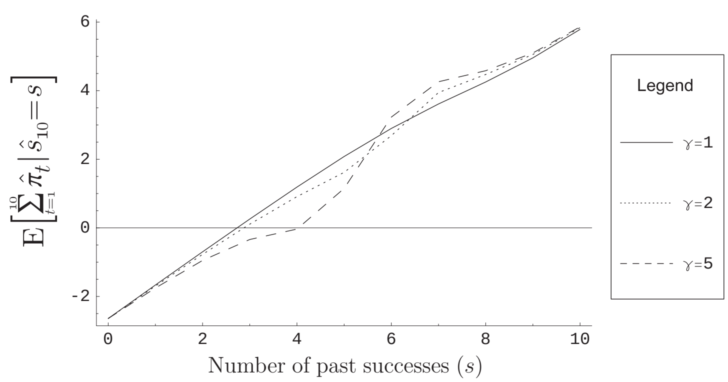

(c) Expected past profits.

Figure 5: Expected insider profits in period 10 (a), expected overconfidence in that period (b), and expected past insider profits (c) as functions the insider's successes in the first 10 periods. The figure was obtained with $H = 0.9$, $L = 0.5$, $\phi_0 = 0.5$, $\Sigma = \Omega = 1$. Each line was obtained with a different value of γ shown in the legend.

that manager's level of overconfidence. If the investor has personal contact with a manager he may try to assess that manager's overconfidence through social cues, but when such cues are not available, our model suggests that a manager's success record will be indicative of his overconfidence.

Using a manager's success record as a measure of his overconfidence creates a dilemma for the investor since the investor uses the same success record to assess the manager's ability. While a trader would always prefer to have as good a success record as possible it is not clear that, when choosing a money manager, an investor will always prefer the one with the best past record. A very successful trader may be too overconfident and therefore trade too aggressively. The effects of overconfidence on trading are likely to be exacerbated when risk aversion and agency issues are introduced to the picture. An overconfident money manager may take risks with his client's money which the client would not endorse. Investors can try to protect themselves from choosing the most overconfident managers by avoiding managers who are successful but underexperienced. They should also judge managers on their long term performance, rather than their most recent successes.[25]

In our model, all insiders, even those with low ability, earn positive expected profits from trading with liquidity traders. In real markets traders who have experienced repeated failures are likely to lose their jobs, their money, or their confidence, and quit trading. The traders who remain will be those with the greatest ability and the greatest overconfidence. This survivorship bias, like the selection bias mentioned in footnote 17, will make the overconfidence level of those active in the marketplace higher than that of the general population. This is in contrast to the results of the natural selection literature,[26] which argue that overconfident and irrational traders will be driven out of financial markets over time. This does not happen here, since trading profits are what make insiders overconfident.

We finish this section by looking at the volatility of prices conditional on the insider's past performance. Expected volatility is not affected by the insider's successes in the same way as are overconfidence and volume. Although expected overconfidence and expected trading volume can both be non-monotonic in the number of past insider successes, expected volatility is always increased by one more insider success. More precisely, large posteriors by the biased insider ($\bar{\mu}_t(s)$) and the rational market maker ($\mu_t(s)$) both contribute to more expected volatility: the former by

[25] As discussed in section 3.2, while we can identify the factors that determine on average how much time it will take for a trader to achieve his maximum overconfidence, we are unable to say how long this time is for any class of traders such as money managers. Furthermore, since our model allows for only two ability levels, there is not much room for a trader who has a long and successful track record to overestimate his ability. That is, he thinks he is of ability type H, and he is probably right. In real markets, traders can always believe themselves to have more ability than they have. Thus even those with long histories of success may fall prey to hubris.

[26] See, for example, Blume and Easley (1982, 1992), and Luo (1998).

his unwarranted aggressiveness, and the latter by his steeper price schedule.[27]

Proposition 3.8 *At the end of period t, the conditional expected volatility in period $t+1$ is increasing in the number of past successful predictions by the insider in the first t periods.*

Proof : See Appendix B.

3.5 Price Patterns

Daniel, Hirshleifer and Subrahmanyam (1998) argue that a trader's overconfidence in his private information can result in positive (negative) price autocorrelation in the short (long) run. In their model, overconfident traders initially overreact to their long-lived private information about a security's payoff, but eventually realize their mistakes through the gradual public revelation of that payoff.

In contrast, the information in our model is short-lived, in the sense that each signal obtained by the insider is advantageous to him for one and only one period. In other words, the insider can profit from each piece of information by trading only once, after which the information is instantaneously made public. As a result, consecutive market-clearing prices are independent in our model: $\text{Cov}(\hat{p}_t, \hat{p}_{t+1}) = 0$, for all $t = 1, 2, \ldots$. Furthermore, this implies that the returns (i.e. the price changes) are spuriously negatively autocorrelated through a bid-ask bounce, an effect originally documented by Roll (1984):

$$\text{Cov}(\hat{p}_{t+1} - \hat{p}_t, \hat{p}_{t+2} - \hat{p}_{t+1}) = -\text{Var}(\hat{p}_{t+1}).$$

We know from the last two sections what effects the insider's learning bias and past performance have on the variance of prices; these effects are thus simply reversed for the return autocovariances.

We can still reconcile our results with those of Daniel, Hirshleifer and Subrahmanyam (1998), and gain some insight into the effects of the insider's learning bias on the price process by looking at the expected evolution of prices for a given dividend size. In an economy where the insider's ability is known ex ante (i.e. $H = L = \mu$), the expected price in any period t given a subsequent

[27]It should be noted that the increase in volatility is the result of the overconfident insider revealing, through aggressive trading, more of his signal to the market maker than is optimal. This allows the market maker to set prices closer to the actual forthcoming dividend, $\hat{v}_t$, and further from its unconditional mean. If, as here, the insider's information is revealed in the next period, this increase in volatility is short-lived and will not even be detected if prices are measured only in the periods following public revelations.

announcement of $\hat{v}_t = v$ is constant at $\mathrm{E}(\hat{p}_t|\hat{v}_t = v) = \frac{v\mu^2}{2}$. As we next show, this is not true when the insider learns about his own ability through time (when $H > L$), in which case

$$\mathrm{E}(\hat{p}_t|\hat{v}_t = v) = \frac{v}{2}\sum_{s=0}^{t-1}\binom{t-1}{s}\left[H^s(1-H)^{t-1-s}\phi_0 + L^s(1-L)^{t-1-s}(1-\phi_0)\right]\bar{\mu}_{t-1}(s)\mu_{t-1}(s).$$

First, when the insider learns his ability without a bias ($\gamma = 1$), the expected price in period t for a given subsequent positive dividend[28] will be increasing (decreasing) when his ability turns out to be high (low). This is not surprising since the insider learns that he should use his information more (less) aggressively over time. Unconditionally, the expected price starts at $\frac{v\mu^2}{2}$ and increases to its long-run value. These effects are illustrated by the dotted lines in Figure 6 for the same set of parameters as in previous figures.

When the insider learns with a bias, the expected price for a given dividend is always higher than when he learns without the bias. Moreover, it is still the case that the expected price increases to the correct/unbiased value when the insider's ability turns out to be high. However, when the insider's ability is low, convergence to the unbiased long-run expected price only occurs for reasonable levels of self-attribution bias (i.e. $\gamma < \gamma^*$); however, if the insider is so biased as to refuse to acknowledge his low ability despite persistent poor performance (i.e. $\gamma \geq \gamma^*$), the expected price stays too high. Unconditionally, the expected price for a given dividend will therefore converge to the unbiased value for reasonable levels of bias (i.e., $\gamma < \gamma^*$). All these effects are illustrated in Figure 6.

Note that in Figure 6(c), even when expected prices converge to their unbiased long-run value, they have a humped-shape pattern similar to that of expected overconfidence in Figure 2, and similar to that found by Daniel, Hirshleifer and Subrahmanyam (1998). The overconfidence developed over time by the insider causes him to push prices too far in the short run but, in the long run, his learning drives prices back to their correct values.

4 Discussion

Our model predicts that overconfident traders will increase their trading volume and thereby lower their expected profits. To the extent that trading is motivated by overconfidence, higher trading will correlate with lower profits. Barber and Odean (1998a) find that this is true for individual investors.

[28] All the effects are reversed for a negative dividend.

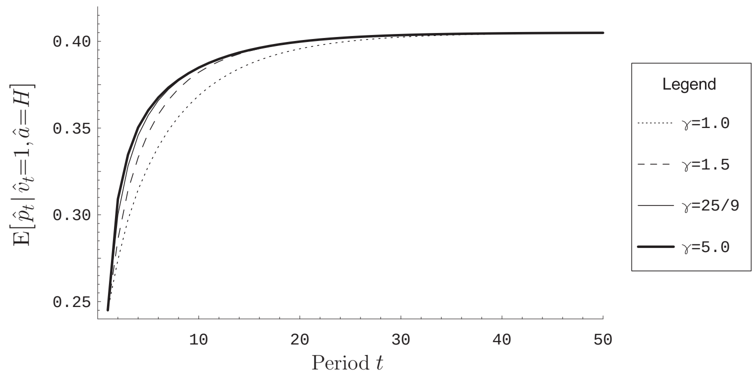

(a) Expected prices when $\hat{a} = H$.

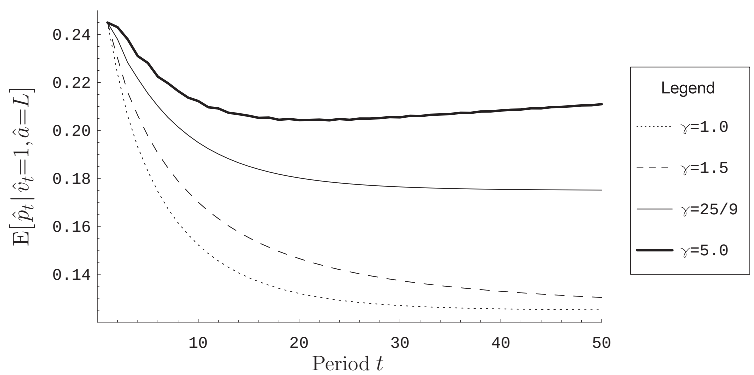

(b) Expected prices when $\hat{a} = L$.

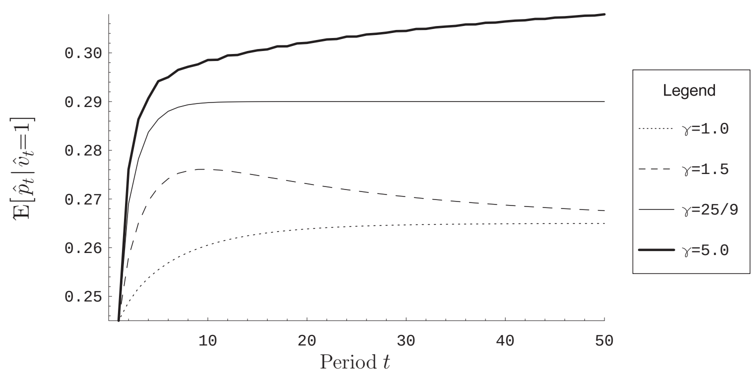

(c) Unconditional expected prices.

Figure 6: Expected patterns of prices for a given dividend size of \$1 per period (a) conditional on $\hat{a} = H$; (b) conditional on $\hat{a} = L$; (c) unconditionally. All three figures were obtained with $H = 0.9$, $L = 0.5$, $\phi_0 = 0.5$, and $\Sigma = \Omega = 1$. Each line was drawn with a different γ, shown in the legends. Note that, with these parameter values, $\gamma^* = 25/9 \approx 2.78$.

While this evidence supports our model, our model does more than simply posit that investors are overconfident. We also describe a dynamic by which overconfidence may wax and wane, both on an individual level and in the aggregate (though the latter is not modeled formally). In times when aggregate success is greater than usual, overconfidence will be higher. In our model, success is measured by how well a trader forecasts dividends. This formulation allows us to present closed form solutions. In many markets, returns will be a trader's metric of success. Traders who attribute returns from general market increases to their own acumen will become overconfident and therefore trade more actively. Therefore, we would predict that periods of market increases will tend to be followed by periods of increased aggregate trading. Statman and Thorley (1998) find this is so for monthly horizons.[29] Taking a longer view, overconfidence and its principal side effect, increased trading, are likely to rise late in a bull market and to fall late in a bear market. A bull market may also attract more investment capital, in part, because investors grow more confident in their personal investment abilities. This increase in investment capital could cause price pressures that sent market prices even higher.[30]

In our model, investors are most overconfident early in their careers. With more experience, self-assessment becomes more realistic and overconfidence more subdued. Barber and Odean (1998b) find that, after controlling for gender, marital status, children, and income, younger investors trade more actively than older investors while earning lower returns relative to a buy-and-hold portfolio. These results are consistent with our prediction that overconfidence diminishes with greater experience.

A further testable empirical prediction of our model is that, on an individual level, investors who realize abnormally good returns in one period will, on average, trade more actively in the next period and in so doing lower their net returns. While the same prediction could be made for professional money managers, their trading decisions will also be influenced by agency considerations that may be difficult to disentangle from overconfidence.

Finally, it is worth noting that if other behavioral biases affect asset prices, then overconfidence is likely to magnify those effects by giving investors the fortitude to act more aggressively on their biased impulses.

[29] We were informed of Statman and Thorley's results after writing this paper.

[30] Of course other factors can lead to similar results. Investors might upwardly revise their estimate of the market's expected return during a bull market and therefore invest more. Or a bull market might benefit from demographically driven changes in aggregate savings.

5 Conclusion

We go through life learning about ourselves as well as the world around us. We assess our own abilities not so much through introspection as by observing our successes and failures. Most of us tend to take too much credit for our own successes. This leads to overconfidence. It is in this way that overconfidence develops in our model. When a trader is successful, he attributes too much of his success to his own ability and revises his beliefs about his ability upward too much. In our model overconfidence is dynamic, changing with successes and failures. Average levels of overconfidence are greatest in those who have been trading for a short time. With more experience, people develop better self assessments.

Since it is through success that traders become overconfident, successful traders, though not necessarily the most successful traders, are most overconfident. These traders are also, as a result of success, wealthy. Overconfidence does not make traders wealthy, but the process of becoming wealthy can make them overconfident. Thus overconfident traders can play an important long-term role in financial markets.

As shown in our model, an overconfident trader trades too aggressively, and this increases expected trading volume. Volatility is increasing in a trader's number of past successes (for a given number of periods). Both volume and volatility increase with the degree of a trader's learning bias. Overconfident traders behave suboptimally, thereby lowering their expected profits. A more successful trader is likely to have more information gathering ability but he may not use his information well. Thus the expected future profits of a more successful trader may actually be lower than those of a less successful trader. Successful traders do tend to be good, but not as good as they think they are.

The principal goal of this paper is to demonstrate that a simple and prevalent bias in evaluating one's own performance is sufficient to create markets in which investors are, on average, overconfident. Unlike models such as De Long et al. (1990), in which biased traders survive by earning greater profits, our model describes a market in which overconfident traders realize, on average, lower profits. Though overconfidence does not lead to greater profits, greater profits do lead to overconfidence. A particular trader's overconfidence will not flourish indefinitely; time and experience gradually rid him of it. However, in a market in which new traders are born every minute, overconfidence will flourish.

Appendix A

Proof of Lemma 2.1

Assume that $P_t(\hat{\omega}_t, \hat{s}_{t-1}) = \lambda_t(\hat{s}_{t-1})\,\hat{\omega}_t$. This means that the insider's expected period t profits, when sending a market order of $\hat{x}_t$ to the market maker, are given by

$$\begin{aligned}
\mathrm{E}_b[\hat{\pi}_t \mid \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t] &= \mathrm{E}_b\left\{\hat{x}_t[\hat{v}_t - P_t(\hat{\omega}_t, \hat{s}_{t-1})] \,\middle|\, \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t\right\} \\
&= \mathrm{E}_b\left\{\hat{x}_t[\hat{v}_t - \lambda_t(\hat{s}_{t-1})\hat{\omega}_t] \,\middle|\, \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t\right\} \\
&= \mathrm{E}_b\left\{\hat{x}_t[\hat{v}_t - \lambda_t(\hat{s}_{t-1})(\hat{x}_t + \hat{z}_t)] \,\middle|\, \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t\right\} \\
&= \hat{x}_t\left[\mathrm{E}_b(\hat{v}_t \mid \hat{\theta}_t, \hat{s}_{t-1}) - \lambda_t(\hat{s}_{t-1})\hat{x}_t\right], \qquad (28)
\end{aligned}$$

where the last equality follows from the fact that $\hat{z}_t$ is independent from both $\hat{\theta}_t$ and $\hat{s}_{t-1}$, and has a mean of zero. Differentiating this last expression with respect to $\hat{x}_t$ and setting the result equal to zero yields

$$\hat{x}_t = \frac{\mathrm{E}_b(\hat{v}_t \mid \hat{\theta}_t, \hat{s}_{t-1})}{2\lambda_t(\hat{s}_{t-1})}. \qquad (29)$$

Also, a simple use of iterated expectations and the projection theorem for normal variables[31] shows that

$$\begin{aligned}
\mathrm{E}_b(\hat{v}_t \mid \hat{\theta}_t, \hat{s}_{t-1}) &= \mathrm{E}_b\left[\mathrm{E}_b(\hat{v}_t|\hat{\theta}_t, \hat{s}_{t-1}, \hat{\delta}_t) \,\middle|\, \hat{\theta}_t, \hat{s}_{t-1}\right] \\
&= \mathrm{E}_b\left[\hat{\delta}_t\hat{\theta}_t + (1-\hat{\delta}_t)\cdot 0 \,\middle|\, \hat{\theta}_t, \hat{s}_{t-1}\right] \\
&= \mathrm{E}_b\left[\hat{\delta}_t \mid \hat{s}_{t-1}\right]\hat{\theta}_t \\
&= \mathrm{E}_b\left[\hat{a} \mid \hat{s}_{t-1}\right]\hat{\theta}_t \\
&= \bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t, \qquad (30)
\end{aligned}$$

[31] The projection theorem for normal variables is as follows: suppose that

$$\begin{bmatrix} \hat{x} \\ \hat{y} \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} \mu_x \\ \mu_y \end{bmatrix}, \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{12} & \Sigma_{22} \end{bmatrix}\right).$$

Then $\mathrm{E}[\hat{x}|\hat{y}] = \mu_x + \frac{\Sigma_{12}}{\Sigma_{22}}(\hat{y} - \mu_y)$.

where the third equality results from the fact that $\hat{\theta}_t$, without $\hat{v}_t$ does not contain any information about $\hat{\delta}_t$ (or, equivalently, about $\hat{a}$).[32] This yields (15).

Next, Assume that $X_t(\hat{\theta}_t, \hat{s}_{t-1}) = \beta_t(\hat{s}_{t-1})\,\hat{\theta}_t$. As discussed in section 1, the market maker's price is a function of the information he gathers from the order flow that is sent to him. More precisely,

$$\begin{aligned}
\hat{p}_t &= \mathrm{E}[\hat{v}_t \mid \hat{\omega}_t, \hat{s}_{t-1}] \\
&= \mathrm{E}\left[\mathrm{E}(\hat{v}_t \mid \hat{\omega}_t, \hat{s}_{t-1}, \hat{\delta}_t) \;\middle|\; \hat{\omega}_t, \hat{s}_{t-1}\right] \\
&= \mathrm{E}\left\{\hat{\delta}_t \mathrm{E}\left[\hat{v}_t \;\middle|\; \hat{\omega}_t = \beta_t(\hat{s}_{t-1})\hat{v}_t + \hat{z}_t, \hat{s}_{t-1}\right] + (1-\hat{\delta}_t)\cdot 0 \;\middle|\; \hat{\omega}_t, \hat{s}_t\right\}. \qquad (31)
\end{aligned}$$

Use of the projection theorem for normal variables shows that

$$\mathrm{E}\left[\hat{v}_t \;\middle|\; \hat{\omega}_t = \beta_t(\hat{s}_{t-1})\hat{v}_t + \hat{z}_t, \hat{s}_{t-1}\right] = \frac{\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t,$$

so that we can rewrite (31) as

$$\begin{aligned}
\hat{p}_t &= \mathrm{E}\left[\hat{\delta}_t \frac{\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t \;\middle|\; \hat{\omega}_t, \hat{s}_{t-1}\right] \\
&= \mathrm{E}[\hat{\delta}_t \mid \hat{s}_{t-1}]\frac{\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t \\
&= \mathrm{E}[\hat{a} \mid \hat{s}_{t-1}]\frac{\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t \\
&= \frac{\mu_{t-1}(\hat{s}_{t-1})\beta_t(\hat{s}_{t-1})\Sigma}{\beta_t^2(\hat{s}_{t-1})\Sigma + \Omega}\hat{\omega}_t, \qquad (32)
\end{aligned}$$

as in (16). ∎

Proof of Lemma 2.2

To see this, recall from (28)-(30) that the insider chooses $\hat{x}_t$ so as to maximize his expected

[32] Again, this is the case since both $\hat{v}_t$ and $\hat{\varepsilon}_t$ have the same distribution.

profits in period t, which can be written as

$$\begin{aligned}\mathrm{E}_b[\hat{\pi}_t \mid \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t] &= \hat{x}_t \left[\mathrm{E}_b(\hat{v}_t \mid \hat{\theta}_t, \hat{s}_{t-1}) - \lambda_t(\hat{s}_{t-1})\hat{x}_t\right] \\ &= \hat{x}_t \left[\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t - \lambda_t(\hat{s}_{t-1})\hat{x}_t\right]. \end{aligned} \tag{33}$$

Assuming that the slope $\lambda_t(\hat{s}_{t-1})$ of the market maker's linear price schedule is positive,[33] we can maximize this expression with respect to $\hat{x}_t$ to obtain the insider's demand:

$$\hat{x}_t = \frac{\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t}{2\lambda_t(\hat{s}_{t-1})}. \tag{34}$$

Of course, this demand is not the same as that of a rational but otherwise identical insider, who instead would be maximizing unbiased expected profits:

$$\mathrm{E}[\hat{\pi}_t \mid \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t] = \hat{x}_t \left[\mu_{t-1}(\hat{s}_{t-1})\hat{\theta}_t - \lambda_t(\hat{s}_{t-1})\hat{x}_t\right]. \tag{35}$$

Since the maket maker is rational in this model, he knows that the (biased) insider's correct expected profits are given by (35), using the suboptimal demand $\hat{x}_t$ as calculated in (34):

$$\begin{aligned}\mathrm{E}[\hat{\pi}_t \mid \hat{\theta}_t, \hat{s}_{t-1}, \hat{x}_t] &= \frac{\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t}{2\lambda_t(\hat{s}_{t-1})} \left[\mu_{t-1}(\hat{s}_{t-1})\hat{\theta}_t - \lambda_t(\hat{s}_{t-1})\frac{\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t}{2\lambda_t(\hat{s}_{t-1})}\right] \\ &= \frac{\bar{\mu}_{t-1}(\hat{s}_{t-1})\hat{\theta}_t^2}{4\lambda_t(\hat{s}_{t-1})}\left[2\mu_{t-1}(\hat{s}_{t-1}) - \bar{\mu}_{t-1}(\hat{s}_{t-1})\right]. \end{aligned} \tag{36}$$

Suppose first that the market maker quotes a price schedule with a positive slope. On average, he then expects to profit from the liquidity trader; in fact, his expected profits against the liquidity trader can be shown to be equal to $\lambda_t(\hat{s}_{t-1})\Omega$. To perform his market clearing duties competitively, it must therefore be the case that the market maker loses that same amount to the insider on average, i.e. it must be the case that (36) is positive. However, when $2\mu_{t-1}(\hat{s}_{t-1}) < \bar{\mu}_{t-1}(\hat{s}_{t-1})$, this is not the case, so that an equilibrium with a positively sloped price schedule is impossible.

What happens if the market maker quotes a price schedule with a negative slope? From (33),

[33] This is needed in order to satisfy the second-order condition. If the slope of the price schedule is negative, then the insider will want to trade an infinite number of shares.

we see that the insider's problem degenerates, as he would then choose an infinite demand. This would not only make his biased expected profits infinite (and positive), but would also make his unbiased expected profits in (35) infinite (and negative). More than that, any negatively sloped price schedule implies that the market maker will also lose against the liquidity trader; his expected losses are given by $-\lambda_t(\hat{s}_{t-1})\Omega$. It is therefore impossible for the market maker to perform his duties competitively with any negatively sloped price schedule. ∎

Proof of Lemma 2.3

The sufficiency part of the argument is obvious as $H \leq 2L$ implies that

$$2\mu_{t-1}(\hat{s}_{t-1}) \geq 2L \geq H \geq \bar{\mu}_{t-1}(\hat{s}_{t-1}).$$

To show necessity, we show that if $2L < H$, then $2\mu_t(s) < \bar{\mu}_t(s)$ for some integers s and t such that $0 < s \leq t$, and some $\gamma > 1$. So, suppose that $2L < H$. For any $\epsilon > 0$, it is possible to find integers s and t such that $0 < s \leq t$ and

$$\mu_t(s) \leq L + \epsilon.$$

Since $\bar{\mu}_t(s)$ increases to H as γ increases to infinity, it is also possible to find $\gamma > 1$ such that

$$\bar{\mu}_t(s) \geq H - \epsilon.$$

By choosing ϵ to be strictly smaller than $\frac{H-2L}{3}$, we have

$$2\mu_t(s) \leq 2(L + \epsilon) < H - \epsilon \leq \bar{\mu}_t(s).$$

This completes the proof. ∎

Proof of Lemma 2.4

By using (18) in (17) and rearranging, we obtain

$$2\mu_{t-1}(s)\Sigma\beta_t^2(s) = \bar{\mu}_{t-1}(s)\Sigma\beta_t^2(s) + \bar{\mu}_{t-1}(s)\Omega,$$

which is quadratic in $\beta_t(s)$. As long as $2\mu_{t-1}(s) \geq \bar{\mu}_{t-1}(s)$, we can solve for $\beta_t(s)$ and obtain (19a), as desired.[34] Also, as argued above, a necessary and sufficient condition for this inequality to be satisfied for any integers s and t such that $0 \leq s < t$ and any $\gamma > 1$ is that $H \leq 2L$. Finally, using (19a) for $\beta_t(s)$ in (18) yields (19b). ∎

[34]The other root is rejected, since it represents a minimum, not a maximum.

Appendix B

Proof of Lemma 3.1

As discussed in the paragraph preceding the lemma, all we need to show is that

$$\left(\frac{L}{\gamma H}\right)^{H}\left(\frac{1-L}{1-H}\right)^{1-H} < 1.$$

By taking log's of both sides and rearranging, this can be shown to be equivalent to showing that

$$f_{\gamma,L}(H) \equiv (1-H)\log(1-L) - (1-H)\log(1-H) - H\log\gamma - H\log H + H\log L < 0 \qquad (37)$$

for all $H \in (L,1]$. First, note that $f_{\gamma,L}(L) = -L\log\gamma \leq 0$. So, if we can show that $f'_{\gamma,L}(H) < 0$ for all $H \in (L,1]$, we will have the desired result. Indeed,

$$\begin{aligned} f'_{\gamma,L}(H) &= -\log(1-L) + \log(1-H) - \log\gamma - \log H + \log L \\ &= -\log\gamma - \log\left(\frac{H(1-L)}{L(1-H)}\right) < 0, \end{aligned}$$

since $\gamma \geq 1$ and $\frac{H(1-L)}{L(1-H)} > 1$. ∎

Proof of Lemma 3.2

As $t \to \infty$, since $\hat{a} = L$, we expect the insider to correctly guess the one-period dividend a fraction L of the times. So, as we play the game more and more often (t tends to ∞), we expect his updated posteriors $\bar{\phi}_t(s)$ to behave like

$$\frac{1}{1+\left(\frac{L}{\gamma H}\right)^{Lt}\left(\frac{1-L}{1-H}\right)^{t-Lt}\frac{1-\phi_0}{\phi_0}} = \frac{1}{1+\left[\left(\frac{L}{\gamma H}\right)^{L}\left(\frac{1-L}{1-H}\right)^{1-L}\right]^{t}\frac{1-\phi_0}{\phi_0}}.$$

So $\bar{\phi}_t(s)$ will converge to 0, ϕ_0, or 1 according to whether the expression in square brackets is greater than, equal to, or smaller than 1. By taking log's, this is equivalent to finding whether

$$g_{\gamma,H}(L) \equiv (1-L)\log(1-L) - (1-L)\log(1-H) - L\log\gamma - L\log H + L\log L$$

is greater than, equal to, or smaller than 0. Let us first check that $g_{1,H}(L) > 0$ for all $L \in [0, H)$, which we do by verifying that $g_{1,H}(H) = 0$, and $g'_{1,H}(L) < 0$ for all $L \in [0, H)$. The first part is easily verified, and

$$\begin{aligned} g'_{1,H}(L) &= -\log(1-L) + \log(1-H) - \log H + \log L \\ &= -\log\left(\frac{H(1-L)}{L(1-H)}\right) < 0, \end{aligned}$$

since $H(1-L) > L(1-H)$. Now, observe that $\frac{\partial}{\partial\gamma} g_{\gamma,H}(L) = -\frac{L}{\gamma} < 0$, so that $g_{1,H}(L) > g_{\gamma,H}(L)$ for all $\gamma \in (1, \infty)$ and $\lim_{\gamma\to\infty} g_{\gamma,H}(L) = -\infty$. Since $g_{1,H}(L) > 0$, this means that there will always be a value γ^* such that

$$g_{\gamma,H}(L) \begin{cases} > 0, & \text{if } \gamma < \gamma^* \\ = 0, & \text{if } \gamma = \gamma^* \\ < 0, & \text{if } \gamma > \gamma^*. \end{cases}$$

This value γ^* solves $g_{\gamma^*,H}(L) = 0$, and it is easily shown to be given by

$$\gamma^* = \frac{L}{H}\left(\frac{1-L}{1-H}\right)^{(1-L)/L}.$$

This completes the proof. ▮

Proof of Proposition 3.1

Since the denominator of $K_t(s)$ in (20) is not a function of γ, $\frac{\partial K_t(s)}{\partial\gamma}$ will have the same sign as $\frac{\partial\bar{\mu}_t(s)}{\partial\gamma}$. We show in Result C.2 of Appendix C that $\frac{\partial\bar{\mu}_t(s)}{\partial\gamma} > 0$. ▮

Proof of Proposition 3.2

In our model, the number of successes in the first t periods is obviously an integer in $\{0, 1, \ldots, t\}$,

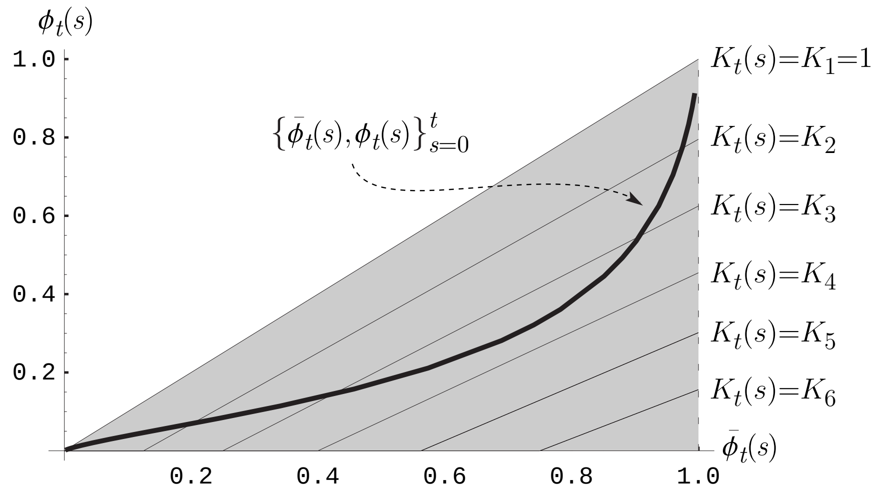

Figure 7: This figure shows $\phi_t(s)$ as a function of $\bar{\phi}_t(s)$. For any $s \in [0, t]$, we have $\bar{\phi}_t(s) \geq \phi_t(s)$, so that all the points $\{\bar{\phi}_t(s), \phi_t(s)\}_{s=0}^t$ must lie in the grey area. The thin solid lines represent the "iso-confidence" curves $K_t(s) = K_i$, $i = 1, \dots, 6$ for $1 = K_1 < K_2 < \cdots < K_6$. The thick solid line represents the parametric curve $\{\bar{\phi}_t(s), \phi_t(s)\}_{s=0}^t$, where $\phi_t(s)$ and $\bar{\phi}_t(s)$ are given in (5) and (8) respectively.

but the function $K_t(s)$ is well defined for any $s \in [0, t]$. We first show that this function is increasing for $s \in [0, s_0]$ and decreasing for $s \in [s_0, t]$ for some $s_0 \in [0, t]$.

To show this, recall that

$$K_t(s) = \frac{\bar{\mu}_t(s)}{\mu_t(s)} = \frac{L + (H - L)\bar{\phi}_t(s)}{L + (K - L)\phi_t(s)}.$$

If we define an "iso-confidence" curve by $K_t(s) = K_i$ for some constant $K_i \geq 1$, each of these curves can then be written as a straight line in a $\bar{\phi}_t(s)$-$\phi_t(s)$ diagram. More precisely, each iso-confidence curve can be expressed as

$$\phi_t(s) = \frac{1}{K_i}\left[\bar{\phi}_t(s) - \frac{(K_i - 1)L}{H - L}\right].$$

These lines are shown as thin solid lines for $1 = K_1 < K_2 < \cdots < K_6$ in Figure 7.

From Result C.1 in Appendix C, we know that the parametric curve $\{\bar{\phi}_t(s), \phi_t(s)\}_{s=0}^t$ starts at $(0, 0)$ and is increasing. This curve is shown as a thick solid line in Figure 7. Since the iso-confidence curves are linear in this $\bar{\phi}_t(s)$-$\phi_t(s)$ diagram, we only need to show that $\phi_t(s)$ in first concave and then convex as a function of $\bar{\phi}_t(s)$. Indeed, it will then be the case that each iso-confidence curve is crossed at most twice or, equivalently, that $K_t(s)$ is increasing and then decreasing as a function of s.

To show this, we use Result C.1 in Appendix C to obtain

$$\frac{\partial \phi_t(s)}{\partial \bar{\phi}_t(s)} = \frac{\partial \phi_t(s)/\partial s}{\partial \bar{\phi}_t(s)/\partial s} = \frac{\phi_t(s)[1-\phi_t(s)] \log\left(\frac{H}{L}\frac{1-L}{1-H}\right)}{\bar{\phi}_t(s)[1-\bar{\phi}_t(s)] \log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right)}, \tag{38}$$

and

$$\frac{\partial}{\partial s}\left(\frac{\partial \phi_t(s)}{\partial \bar{\phi}_t(s)}\right) = \tag{39}$$
$$\frac{\log\left(\frac{H}{L}\frac{1-L}{1-H}\right)}{\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right)} \frac{\phi_t(s)[1-\phi_t(s)]}{\bar{\phi}_t(s)[1-\bar{\phi}_t(s)]} \left\{[1-2\phi_t(s)] \log\left(\frac{H}{L}\frac{1-L}{1-H}\right) - [1-2\bar{\phi}_t(s)] \log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right)\right\}.$$

Using standard calculus results along with (39) and Result C.1 in Appendix C, we have

$$\frac{\partial^2 \phi_t(s)}{\partial \bar{\phi}_t(s)^2} = \frac{\frac{\partial}{\partial s}\left(\frac{\partial \phi_t(s)}{\partial \bar{\phi}_t(s)}\right)}{\partial \bar{\phi}_t(s)/\partial s} =$$
$$\frac{\phi_t(s)[1-\phi_t(s)]}{\left\{\bar{\phi}_t(s)[1-\bar{\phi}_t(s)]\right\}^2} \frac{[1-2\phi_t(s)] \log\left(\frac{H}{L}\frac{1-L}{1-H}\right) - [1-2\bar{\phi}_t(s)] \log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right)}{\log\left(\frac{H}{L}\frac{1-L}{1-H}\right)}.$$

This last expression will always have the same sign as

$$D\left(\phi_t(s), \bar{\phi}_t(s)\right) \equiv [1-2\phi_t(s)] \log\left(\frac{H}{L}\frac{1-L}{1-H}\right) - [1-2\bar{\phi}_t(s)] \log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right). \tag{40}$$

Since $\gamma \geq 1$, (40) is negative for $(\bar{\phi}_t(s), \phi_t(s)) = (0, 0)$ and, since $\bar{\phi}_t(s) \geq \phi_t(s)$, it is positive for $\bar{\phi}_t(s) \geq 1/2$. Therefore, if we can show that $D\left(\phi_t(s), \bar{\phi}_t(s)\right)$ is increasing in $\bar{\phi}_t(s)$ for $\bar{\phi}_t(s) \in [0, 1/2]$, we will have the desired result. Using (38), we have

$$\begin{aligned}
\frac{dD\left(\phi_t(s),\bar{\phi}_t(s)\right)}{d\bar{\phi}_t(s)} &= \frac{\partial D\left(\phi_t(s),\bar{\phi}_t(s)\right)}{\partial \phi_t(s)}\frac{\partial \phi_t(s)}{\partial \bar{\phi}_t(s)} + \frac{\partial D\left(\phi_t(s),\bar{\phi}_t(s)\right)}{\partial \bar{\phi}_t(s)} \\
&= -2\log\left(\frac{H}{L}\frac{1-L}{1-H}\right)\frac{\phi_t(s)[1-\phi_t(s)]\log\left(\frac{H}{L}\frac{1-L}{1-H}\right)}{\bar{\phi}_t(s)[1-\bar{\phi}_t(s)]\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right)} + 2\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right).
\end{aligned}$$

This last expression is nonnegative if and only if

$$\bar{\phi}_t(s)[1-\bar{\phi}_t(s)]\log^2\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right) \geq \phi_t(s)[1-\phi_t(s)]\log^2\left(\frac{H}{L}\frac{1-L}{1-H}\right).$$

Since $\gamma > 1$, this is definitely true for $0 \leq \phi_t(s) \leq \bar{\phi}_t(s) \leq 1/2$. This shows that $K_t(s)$ is first increasing and then decreasing as a function of s.

To complete the proof, we must deal with the fact that, in our problem, we only care about $K_t(s)$ for $s \in \{0,1,\ldots,t\}$. However, the result is immediate from the shape of $K_t(s)$. ∎

Proof of Lemma 3.3

First, a standard result for normal variables is that, if $\hat{y} \sim \mathrm{N}(0,\sigma^2)$, then

$$\mathrm{E}\,|\hat{y}| = \sqrt{\frac{2\sigma^2}{\pi}}.$$

We can therefore write

$$\begin{aligned}
\mathrm{E}\left[\hat{\psi}_{t+1}\,\middle|\,\hat{s}_t = s\right] &= \frac{1}{2}\mathrm{E}\left[|\hat{x}_{t+1}| + |\hat{z}_{t+1}|\,\middle|\,\hat{s}_t = s\right] \\
&= \frac{1}{2}\mathrm{E}\left[|\hat{x}_{t+1}|\,\middle|\,\hat{s}_t = s\right] + \frac{1}{2}\sqrt{\frac{2\Omega}{\pi}} \\
&= \frac{1}{2}\beta_{t+1}(s)\,\mathrm{E}\left[|\hat{\theta}_{t+1}|\,\middle|\,\hat{s}_t = s\right] + \sqrt{\frac{\Omega}{2\pi}} \\
&= \frac{1}{2}\beta_{t+1}(s)\sqrt{\frac{2\Sigma}{\pi}} + \sqrt{\frac{\Omega}{2\pi}} \\
&= \frac{1}{\sqrt{2\pi}}\left[\beta_{t+1}(s)\sqrt{\Sigma} + \sqrt{\Omega}\right],
\end{aligned}$$

and this last expression is equal to (24). The expression for expected profits is derived as follows:

$$
\begin{aligned}
&\mathrm{E}\left[\hat{\pi}_{t+1} \mid \hat{s}_t = s\right] \\
&\quad = \mathrm{E}\left[\hat{x}_{t+1}(\hat{v}_{t+1} - \hat{p}_{t+1}) \mid \hat{s}_t = s\right] \\
&\quad = \mathrm{E}\left(\beta_{t+1}(\hat{s}_t)\hat{\theta}_{t+1}\left\{\hat{v}_{t+1} - \lambda_{t+1}(\hat{s}_t)\left[\beta_{t+1}(\hat{s}_t)\hat{\theta}_{t+1} + \hat{z}_{t+1}\right]\right\} \;\Big|\; \hat{s}_t = s\right) \\
&\quad = \mathrm{E}\left[\mathrm{E}\left(\beta_{t+1}(\hat{s}_t)\hat{\theta}_{t+1}\left\{\hat{v}_{t+1} - \lambda_{t+1}(\hat{s}_t)\left[\beta_{t+1}(\hat{s}_t)\hat{\theta}_{t+1} + \hat{z}_{t+1}\right]\right\} \;\Big|\; \hat{\theta}_{t+1}, \hat{s}_t = s\right) \;\Big|\; \hat{s}_t = s\right] \\
&\quad = \mathrm{E}\left\{\beta_{t+1}(s)\hat{\theta}_{t+1}\left[\mu_t(s)\hat{\theta}_{t+1} - \lambda_{t+1}(s)\beta_{t+1}(s)\hat{\theta}_{t+1}\right] \;\Big|\; \hat{s}_t = s\right\} \\
&\quad = \beta_{t+1}(s)\left[\mu_t(s) - \lambda_{t+1}(s)\beta_{t+1}(s)\right]\mathrm{E}\left(\hat{\theta}^2_{t+1} \mid \hat{s}_t = s\right) \\
&\quad = \beta_{t+1}(s)\left[\mu_t(s) - \lambda_{t+1}(s)\beta_{t+1}(s)\right]\Sigma.
\end{aligned}
$$

Finally, using the expressions derived for $\beta_t(s)$ and $\lambda_t(s)$ in Lemma 2.4, it is easy to show that

$$\beta_{t+1}(s)\left[\mu_t(s) - \lambda_{t+1}(s)\beta_{t+1}(s)\right] = \frac{1}{2}\sqrt{\frac{\Omega}{\Sigma}\bar{\mu}_t(s)\left[2\mu_t(s) - \bar{\mu}_t(s)\right]},$$

so that

$$\mathrm{E}\left[\hat{\pi}_{t+1} \mid \hat{s}_t = s\right] = \frac{1}{2}\sqrt{\Sigma\Omega}\,\sqrt{\bar{\mu}_t(s)\left[2\mu_t(s) - \bar{\mu}_t(s)\right]},$$

as in (25). In this economy, the security price in a period reflects only the dividend paid at the end of that period. Since the dividend's unconditional mean is zero, the price's unconditional mean is also zero. Therefore,

$$
\begin{aligned}
&\mathrm{Var}(\hat{p}_{t+1} \mid \hat{s}_t = s) \\
&\quad = \mathrm{E}\left(\hat{p}^2_{t+1} \mid \hat{s}_t = s\right) \\
&\quad = \mathrm{E}\left[\lambda^2_{t+1}(\hat{s}_t)\,\hat{\omega}^2_{t+1} \mid \hat{s}_t = s\right] \\
&\quad = \mathrm{E}\left[\lambda^2_{t+1}(\hat{s}_t)(\hat{x}_{t+1} + \hat{z}_{t+1})^2 \mid \hat{s}_t = s\right] \\
&\quad = \mathrm{E}\left\{\lambda^2_{t+1}(\hat{s}_t)\left[\beta_{t+1}(\hat{s}_t)\hat{\theta}_{t+1} + \hat{z}_{t+1}\right]^2 \;\Big|\; \hat{s}_t = s\right\} \\
&\quad = \lambda^2_{t+1}(s)\left[\beta^2_{t+1}(s)\mathrm{E}(\hat{\theta}^2_{t+1} \mid \hat{s}_t = s) + 2\beta_{t+1}(s)\mathrm{E}(\hat{\theta}_{t+1}\hat{z}_{t+1} \mid \hat{s}_t = s) + \mathrm{E}(\hat{z}^2_{t+1} \mid \hat{s}_t = s)\right] \\
&\quad = \lambda^2_{t+1}(s)\left[\beta^2_{t+1}(s)\Sigma + \Omega\right].
\end{aligned}
$$

Now, using the expressions derived for $\beta_t(s)$ and $\lambda_t(s)$ in Lemma 2.4, it is easy to show that

$$\lambda_{t+1}^2(s)\left[\beta_{t+1}^2(s)\Sigma+\Omega\right]=\frac{\Sigma}{2}\bar{\mu}_t(s)\mu_t(s),$$

so that

$$\mathrm{Var}\left(\hat{p}_{t+1}\mid\hat{s}_t=s\right)=\frac{\Sigma}{2}\,\bar{\mu}_t(s)\,\mu_t(s),$$

as in (26). ∎

Proof of Proposition 3.3

Given the expression for the conditional expected volume in (24), it is sufficient to prove that

$$\frac{\partial\beta_t(s)}{\partial\gamma}>0.$$

Straighforward differentiation of the expression for $\beta_t(s)$ in equation (19a) of Lemma 2.4 results in

$$\frac{\partial\beta_t(s)}{\partial\gamma}=\sqrt{\frac{\Omega}{\Sigma}\frac{2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)}{\bar{\mu}_{t-1}(s)}}\frac{\mu_{t-1}(s)}{\left[2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)\right]^2}\frac{\partial\bar{\mu}_{t-1}(s)}{\partial\gamma},$$

which in turn shows that it is sufficient to show that

$$\frac{\partial\bar{\mu}_{t-1}(s)}{\partial\gamma}>0.$$

This is shown to be true in Result C.2 of Appendix C. ∎

Proof of Proposition 3.4

To show the desired result, we only need to show that $\bar{\mu}_{t-1}(s)\left[2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)\right]$ is decreasing in γ. This is straightforward to show since

$$\frac{\partial}{\partial\gamma}\left\{\bar{\mu}_{t-1}(s)\left[2\mu_{t-1}(s)-\bar{\mu}_{t-1}(s)\right]\right\}=-2\left[\bar{\mu}_{t-1}(s)-\mu_{t-1}(s)\right]\frac{\partial}{\partial\gamma}\bar{\mu}_{t-1}(s),$$

and $\frac{\partial}{\partial\gamma}\bar{\mu}_{t-1}(s)$ is shown to be positive in Result C.2 of Appendix C. ∎

Proof of Proposition 3.5

The result easily follows from the fact that $\frac{\partial}{\partial\gamma}\bar{\mu}_{t-1}(s) > 0$, which is shown to be true in Result C.2 of Appendix C. ∎

Proof of Proposition 3.6

As shown in Lemma 3.3, the expected volume in period $t+1$ is proportional to the expected trading intensity $\beta_{t+1}(s)$ of the insider in that period. Since $\beta_t(s)$ is monotonically increasing in $K_t(s)$ (see equation (27)), the result of Proposition 3.2 immediately implies this result. ∎

Proof of Proposition 3.7

This result is shown in essentially the same way that Proposition 3.2 was shown earlier, except that the "iso-profit" curves are now quadratic. ∎

Proof of Proposition 3.8

In view of (26), this amounts to showing that the product $\bar{\mu}_t(s)\mu_t(s)$ is increasing in s. However, since both these quantities are increasing in s (see Result C.2 in Appendix C), the result follows easily. ∎

Appendix C

This appendix contains a few results that are used in the proofs of some propositions in section 3.

Result C.1 *The partial derivatives of* $\bar{\phi}_t(s)$ *in (8) with respect to* γ *and* s *are respectively equal to*

$$\frac{\partial \bar{\phi}_t(s)}{\partial \gamma} = \frac{s}{\gamma}\bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right] \geq 0, \tag{41}$$

and

$$\frac{\partial \bar{\phi}_t(s)}{\partial s} = \bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right]\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right) \geq 0. \tag{42}$$

Proof: Staightforward differentiation of $\bar{\phi}_t(s)$ yields

$$\begin{aligned}
\frac{\partial \bar{\phi}_t(s)}{\partial \gamma} &= \Big\{ s\gamma^{s-1}H^s(1-H)^{t-s}\phi_0\left[(\gamma H)^s(1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)\right] - \\
&\quad (\gamma H)^s(1-H)^{t-s}\phi_0 s\gamma^{s-1}H^s(1-H)^{t-s}\phi_0 \Big\} \\
&\qquad \div \left[(\gamma H)^s(1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)\right]^2 \\
&= \frac{\frac{s}{\gamma}(\gamma H)^s(1-H)^{t-s}\phi_0 L^s(1-L)^{t-s}(1-\phi_0)}{\left[(\gamma H)^s(1-H)^{t-s}\phi_0 + L^s(1-L)^{t-s}(1-\phi_0)\right]^2} \\
&= \frac{s}{\gamma}\bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right],
\end{aligned}$$

which is obviously greater than or equal to zero. Now, since we can write

$$\bar{\phi}_t(s) = \frac{1}{1+\left(\frac{L}{\gamma H}\frac{1-H}{1-L}\right)^s\left(\frac{1-L}{1-H}\right)^t\frac{1-\phi_0}{\phi_0}},$$

we have

$$
\begin{aligned}
\frac{\partial \bar{\phi}_t(s)}{\partial s} &= \frac{(-1)\left(\frac{L}{\gamma H}\frac{1-H}{1-L}\right)^s \log\left(\frac{L}{\gamma H}\frac{1-H}{1-L}\right)\left(\frac{1-L}{1-H}\right)^t \frac{1-\phi_0}{\phi_0}}{\left[1+\left(\frac{L}{\gamma H}\frac{1-H}{1-L}\right)^s\left(\frac{1-L}{1-H}\right)^t\frac{1-\phi_0}{\phi_0}\right]^2} \\
&= \bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right]\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right).
\end{aligned}
$$

Since $\gamma H > L$ and $1-L > 1-H$, this last quantity is obviously greater than or equal to zero. ∎

Result C.2 *The partial derivatives of $\bar{\mu}_t(s)$ in (9) with respect to γ and s are respectively equal to*

$$
\frac{\partial \bar{\mu}_t(s)}{\partial \gamma} = (H-L)\frac{\partial \bar{\phi}_t(s)}{\partial \gamma} = (H-L)\frac{s}{\gamma}\bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right] \geq 0, \tag{43}
$$

and

$$
\frac{\partial \bar{\mu}_t(s)}{\partial s} = (H-L)\frac{\partial \bar{\phi}_t(s)}{\partial s} = (H-L)\bar{\phi}_t(s)\left[1-\bar{\phi}_t(s)\right]\log\left(\frac{\gamma H}{L}\frac{1-L}{1-H}\right) \geq 0. \tag{44}
$$

Proof: Since we have

$$
\bar{\mu}_t(s) = H\bar{\phi}_t(s) + L\left[1-\bar{\phi}_t(s)\right] = L + (H-L)\bar{\phi}_t(s)
$$

and $H > L$, this result follows immediately from Result C.1 above. ∎

Appendix D

In this appendix, we present a few alternative specifications for the model. These specifications show that the results of section 3 are robust to different trading crowds and market clearing mechanisms. The reason why the results of the model are unaffected by these changes is because the learning by the insider(s) is independent from the trading process. More precisely, an insider's strategy in each period depends directly on who he trades with and on the market clearing mechanism. However, once the end-of-period dividend is announced, this insider updates his beliefs based on that dividend, not based on that period's trading. In the interest of space, the formal derivations of these alternative models are not included here; however, the derivations are available from the authors upon request.

D.1 Hedgers

The model of section 1 assumes that some traders will trade randomly for purely exogenous reasons. Of course, since the market maker in that model makes zero profits on average, these liquidity traders fuel the insider's profits. In this section, we replace these liquidity traders by rational traders who have a hedging motive for trading. The specification we adopt is similar to that of Glosten (1989).

Suppose that, at the beginning of every period t, a trader (whom we will refer to as the *hedger*) receives a random endowment $\hat{z}_t$ of the risky security. Suppose also that this trader has negative exponential utility $U(\pi_t) = -e^{-R\pi_t}$ with respect to his profits $\hat{\pi}_t^h$ in period t. It can be shown that, if $(\hat{v}_t, \hat{\varepsilon}_t, \hat{z}_t)'$ is distributed as in (13), the demand of this hedger in period t will be proportional to $\hat{z}_t$.[35]

The rest of the model is solved as before, that is the market maker absorbs the total order flow coming from both the insider and the hedger.[36] It can be shown that the equilibrium will be achieved as long as the hedger's risk aversion coefficient R and the variances Σ and Ω are large enough. Of course, the insider's updating is exactly the same as in our original model, and all the results of section 3 obtain.

[35] Of course, given that this hedger is rational, the proportionality constant is decreasing in the slope of the market maker's price schedule.

[36] In fact, in this model, we can even consider a monopolist market maker who seeks to maximize his expected profits, as in Glosten (1989). This extra layer of generality obtains from the elasticity of the hedger's demand, as opposed to the liquidity trader's inelastic demand.

D.2 A Rational Expectations Economy

The presence of the market maker is also not essential to our model. Instead, the market clearing price in each period could be obtained from a rational expectations equilibrium similar to that of Grossman (1976), and Grossman and Stiglitz (1980).

Suppose that the economy consists of two risk-averse traders with negative exponential utility function $U(\pi_t) = -e^{-R\pi_t}$ with respect to their period t profits $\hat{\pi}_t^1$ and $\hat{\pi}_t^2$. Suppose that the first trader, the insider, is endowed with the private information and learning bias described in section 1, but that the second trader has no private information. Of course, this second trader can infer some of the insider's information through the market clearing price. As in Grossman and Stiglitz (1980), we assume that the risky asset supply in period t is a random variable $\hat{z}_t$.

If $(\hat{v}_t, \hat{\varepsilon}_t, \hat{z}_t)'$ is distributed as in (13), it can be shown that the market clearing price in period t for this economy is given by

$$\hat{p}_t = a_t(\hat{s}_{t-1})\hat{\theta}_t - b_t(\hat{s}_{t-1})\hat{z}_t,$$

for some positive functions $a_t(\cdot)$ and $b_t(\cdot)$. Moreover, the demand $\hat{x}_t$ for the risky asset by the insider is linear in the signal $\hat{\theta}_t$ and the price $\hat{p}_t$, and the demand $\hat{y}_t$ for the risky asset by the uninformed trader is proportional to $\hat{p}_t$:

$$\begin{aligned} \hat{x}_t &= \beta_t(\hat{s}_{t-1})\hat{\theta}_t - \alpha_t(\hat{s}_{t-1})\hat{p}_t; \\ \hat{y}_t &= -\eta_t(\hat{s}_{t-1})\hat{p}_t, \end{aligned}$$

where $\beta_t(\cdot)$, $\alpha_t(\cdot)$, and $\eta_t(\cdot)$ are positive functions. As in Grossman (1976), these traders do not trade when they are both rational, that is as long as the insider remains unsuccessful. However, as discussed in section 3.2, the insider's first successful prediction makes him overconfident for the rest of his existence. Starting then, the two traders will start trading nonzero quantities of the risky asset. In fact, it can even be shown that they would do so without a noisy supply of the risky asset, as they agree to disagree. Also, since the insider's updating is unaffected by this alternative market-clearing mechanism, all the results from section 3 continue to hold.

D.3 Multiple Insiders

A potential criticism of the paper's main model is the fact that only one trader possesses information about the risky security's dividend process. For example, our results in sections 3.1 and 3.2 show that a (moderately) biased trader will on average be overconfident early in his career, and well calibrated later on. Can we conclude that an economy's overconfident traders will consist mainly of young traders? Although it is tempting to answer this question in the affirmative from the results derived so far, our one-insider model only addresses this issue indirectly. In this section, we construct a multi-insider version of the model, which will allow us to tackle this question more directly.

Suppose that a new insider is born every period t with an ability $\hat{a}_t$, and lives for T periods. During period t, these insiders are indexed by their age $\tau = 0, \ldots, T-1$ at the beginning of the period. Suppose also that every period consists of N trading rounds (indexed by n), each of which coincides with a dividend payment by the risky security. As in our main model, we assume that a liquidity trader participates in each such trading round, and that orders are cleared through a competitive market maker. This means that, in every period, T insiders, a liquidity trader and the market maker trade a risky security N times.

To simplify the analysis without affecting the results, we assume that at most one of the insiders receives valuable information in each trading round. This trader is denoted by $\hat{\jmath}_{tn}$:

$$\hat{\jmath}_{tn} \mid (\hat{a}_t, \hat{a}_{t-1}, \ldots, \hat{a}_{t-T+1}) = \begin{cases} 0, & \text{prob. } \frac{\hat{a}_t}{T} \\ 1, & \text{prob. } \frac{\hat{a}_{t-1}}{T} \\ \vdots & \\ T-1, & \text{prob. } \frac{\hat{a}_{t-T+1}}{T} \\ T, & \text{prob. } 1 - \frac{1}{T}\sum_{\tau=0}^{T-1} \hat{a}_{t-\tau}. \end{cases}$$

The signal $\hat{\theta}^{\tau}_{tn}$ received by each trader is given by

$$\hat{\theta}^{\tau}_{tn} = \mathbf{1}_{\{\hat{\jmath}_{tn}=\tau\}} \hat{v}_{tn} + \left(1 - \mathbf{1}_{\{\hat{\jmath}_{tn}=\tau\}}\right) \hat{\varepsilon}^{\tau}_{tn},$$

where $\hat{v}_{tn}$ is the risky security's dividend at the end of the trading round, and $\hat{\varepsilon}^{\tau}_{tn}$, $\tau = 0, 1, \ldots, T-1$, are pure noises. The rest of the model is similar to our main one-insider model in that, at the beginning of every trading round n in period t, the T insiders and the liquidity trader send their demands $\hat{x}_{tn}^{\tau}$, $\tau = 0, 1, \ldots, T-1$, and $\hat{z}_{tn}$ for the risky security, and the market maker absorbs the order flow $\hat{\omega}_{tn} = \sum_{\tau=0}^{T-1} \hat{x}_{tn}^{\tau} + \hat{z}_{tn}$ at a competitive price

$$\hat{p}_{tn} = \mathrm{E}\left[\hat{v}_{tn} \mid \hat{\omega}_{tn}, \hat{s}_{t,n-1}^{0}, \hat{s}_{t,n-1}^{1}, \ldots, \hat{s}_{t,n-1}^{T-1}\right],$$

where $\hat{s}_{t,n-1}^{\tau}$ denotes the number of past successful predictions for insider τ (with $\hat{s}_{t,0}^{\tau} = \hat{s}_{t-1,N}^{\tau-1}$).

If we assume that $\hat{v}_{tn}, \hat{\varepsilon}_{tn}^{0}, \hat{\varepsilon}_{tn}^{1}, \ldots, \hat{\varepsilon}_{tn}^{T-1}$ all have independent normal distribution with a mean of zero and a variance of Σ, and that $\hat{z}_{tn}$ is independently normally distributed with a mean of zero and a variance of Ω, a linear equilibrium exists in each trading round. This linear equilibrium consists in every insider sending a demand proportional to their signal,

$$\hat{x}_{tn}^{\tau} = \beta_{tn}^{\tau}\left(\hat{s}_{t,n-1}^{0}, \hat{s}_{t,n-1}^{1}, \ldots, \hat{s}_{t,n-1}^{T-1}\right)\hat{\theta}_{tn}^{\tau},$$

and the market maker quoting a price schedule proportional to the order flow,

$$\hat{p}_{tn} = \lambda_{tn}\left(\hat{s}_{t,n-1}^{0}, \hat{s}_{t,n-1}^{1}, \ldots, \hat{s}_{t,n-1}^{T-1}\right)\hat{\omega}_{tn}.$$

Unfortunately, closed form solutions are not easily obtainable for this multi-insider model. As a result, we generally have to solve for $\lambda_{tn}\left(s^{0}, s^{1}, \ldots, s^{T-1}\right)$ and $\beta_{tn}^{\tau}\left(s^{0}, s^{1}, \ldots, s^{T-1}\right)$, $\tau = 0, \ldots, T-1$, numerically. Such numerical solutions reveal that all the properties about overconfidence, trading volume, insider profits, and price volatility still hold.

In addition to these results however, we can now tackle the issue of overconfidence as a function of age. This is shown in Figure 8 for a set of parameters similar to that used in Figure 2. This latter figure shows that an insider's overconfidence on average goes up and then down during his lifetime. Figure 8 shows that, in a particular period t, it is the case that the most overconfident traders are relatively young, and that older traders are better calibrated.

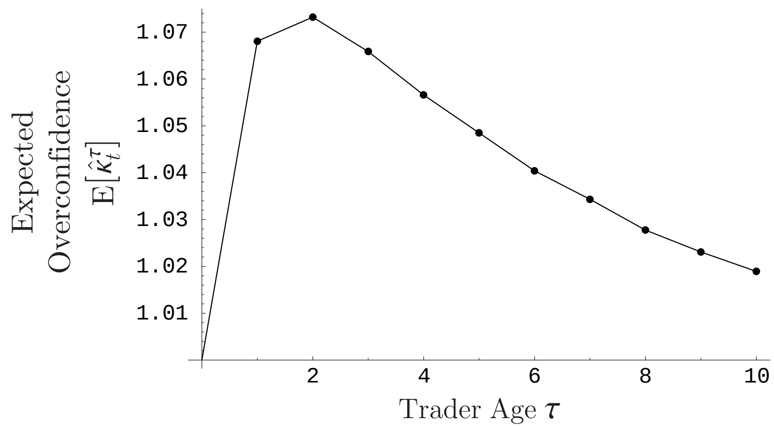

Figure 8: Expected insider overconfidence as a function of the insider's age. In obtaining this figure, every insider's life is set to 10 periods ($T = 10$), and each period consists of 50 trading rounds ($N = 50$). Also, we use the following parameters for all insiders: $H = 0.9$, $L = 0.5$, $\phi_0 = 0.5$, and $\gamma = 1.5$. Finally, we use variances of $\Sigma = \Omega = 1$.

References

Alchian, A. A., 1950, "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, 58, 211-221.

Alpert, M., and H. Raiffa, 1982, "A Progress Report on the Training of Probability Assessors," in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge and New York, 294-305.

Anderson, N., 1959, "Test of a Model for Opinion Change," *Journal of Abnormal and Social Psychology*, 59, 371-381.

Anderson, N., 1981, *Foundations of Information Integration Theory*, Academic Press, New York.

Barber, B., and T. Odean, 1998a, "The Common Stock Investment Performance of Individual Investors," working paper, University of California, Davis.

Barber, B., and T. Odean, 1998b, "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment," working paper, University of California, Davis.

Benos, A., 1998, "Aggressiveness and Survival of Overconfident Traders," *Journal of Financial Markets*, 1, 353-383.

Bhattacharya, U., and M. Spiegel, 1991, "Insiders, Outsiders, and Market Breakdowns," *Review of Financial Studies*, 4, 255-282.

Blume, L., and D. Easley, 1982, "Learning To Be Rational," *Journal of Economic Theory*, 26, 340-351.

Blume, L., and D. Easley, 1992, "Evolution and Market Behavior," *Journal of Economic Theory*, 58, 9-40.

Blume, L., D. Easley, and M. O'Hara, 1994, "Market Statistics and Technical Analysis: The Role of Volume," *Journal of Finance*, 49, 153-181.

Bossaerts, P., 1999, "The Dynamics of Limited-Liability Securities Prices in Fallible Markets," working paper, California Institute of Technology.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam, 1998, "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 53, 1839-1885.

De Finetti, B., 1937, “La prévision: ses lois logiques, ses sources subjectives,” *Annals of the Institut Henri Poincaré*, 7, 1-68.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann, 1990, “Noise Trader Risk in Financial Markets,” *Journal of Political Economy*, 98, 703-738.

Fischhoff, B., 1982, “For Those Condemned to Study the Past: Heuristics and Biases in Hindsight.” In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge and New York, 335-351.

Fiske, S., and S. Taylor, 1991, *Social Cognition* (2nd ed.), McGraw-Hill, New York.

Friedman, M., 1953, *Essays in Positive Economics*, University of Chicago Press, Chicago.

Gervais, S., 1996, “Market Microstructure with Uncertain Information Precision: A Multi-Period Analysis,” working paper, Wharton School, University of Pennsylvania.

Glosten, L., 1989, “Insider Trading, Liquidity, and the Role of the Monopolist Specialist,” *Journal of Business*, 62, 211-235.

Grossman, S. J., 1976, “On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information,” *Journal of Finance*, 31, 573-585.

Grossman, S. J., and J. E. Stiglitz, 1980, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70, 393-408.

Hastorf, A. H., D. J. Schneider, and J. Polefka, 1970, *Person Perception*, Addison-Wesley, Reading, Massachusetts.

Hogarth, R., and H. Einhorn, 1992, “Order Effects in Belief Updating: The Belief-Adjustment Model,” *Cognitive Psychology*, 24, 1-55.

Kyle, A., 1985, “Continuous Auctions and Insider Trading,” *Econometrica*, 53, 1315-1335.

Kyle, A., and F. A. Wang, 1997, “Speculation Duopoly With Agreement to Disagree: Can Overconfidence Survive the Market Test?,” *Journal of Finance*, 52, 2073-2090.

Langer, E., and J. Roth, 1975, “Heads I Win, Tails It’s Chance: The Illusion of Control as a Function of the Sequence of Outcomes in a Purely Chance Task,” *Journal of Personality and Social Psychology*, 32, 951-955.

Lichtenstein, S., B. Fischhoff, and L. Phillips, 1982, "Calibration of Probabilities: The State of the Art to 1980," in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge and New York, 306-334.

Luo, G. Y., 1998, "Market Efficiency and Natural Selection in a Commodity Futures Market," *Review of Financial Studies*, 11, 647-674.

Milgrom, P., and N. Stokey, 1982, "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17-27.

Miller, D., and M. Ross, 1975, "Self-serving Biases in Attribution of Causality: Fact or Fiction?," *Psychological Bulletin*, 82, 213-225.

Odean, T., 1998a, "Volume, Volatility, Price, and Profit When All Traders Are Above Average," *Journal of Finance*, 53, 1887-1934.

Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 39, 1127-1139.

Statman, M., and S. Thorley, 1998, "Overconfidence and Trading Volume," working paper, Santa Clara University.

Subramanyam, K. R., 1996, "Uncertain Precision and Price Reactions to Information," *The Accounting Review*, 71, 207-220.

Wang, A., 1997, "Overconfidence, Delegated Fund Management, and Survival," working paper, Columbia University.

Wolosin, R. J., S. J. Sherman, and A. Till, 1973, "Effects of Cooperation and Competition on Responsibility Attribution After Success and Failure," *Journal of Experimental Social Psychology*, 9, 220-235.