PUNDITS AND QUACKS

By

Jesper Rudiger and Adrien Vigier

March 2015

COWLES FOUNDATION DISCUSSION PAPER NO. 1997

Pundits and Quacks[*]

Jesper Rudiger[†] and Adrien Vigier[‡]

March 16, 2015

Abstract

Do asset prices aggregate investors' private information about the ability of financial analysts? We show that as financial analysts become reputable, the market can get trapped: Investors optimally choose to ignore their private information, and blindly follow analyst recommendations. As time goes by and recommendations accumulate, arbitrage based on the inferred ability of analysts may become profitable again. The market can thus be trapped at times and yet be able, in the long run, to sort the pundits from the quacks. However, this process is impaired when asset fundamentals are volatile: in this case, the market might be trapped indefinitely.

Keywords: Financial Analysts; Reputation; Market Microstructure; Social Learning
JEL Codes: D82; D83; D84; G14; G20

[†]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, DK-1353 Copenhagen, Denmark. *Telephone*: +45 3533 7776. *E-mail*: jesper.rudiger@econ.ku.dk

[‡]Department of Economics, University of Oslo, Moltke Moes Vei 31, N-0851 Oslo, Norway. *Telephone*: +47 2285 4520. *E-mail*: a.h.vigier@econ.uio.no (corresponding author)

1 Introduction

The intensified division of labour taking place within the finance industry raises new and important challenges. Sell-side and independent financial analysts now account for a major part of the information used by investors to determine what stocks to trade.[1] According to Ramnath et al. (2008): "Analysts are viewed as sophisticated processors of financial information who are less likely than naive investors to misunderstand the implications of financial information". Yet while independent analysts have the potential to improve markets' efficiency, their existence paves the way for a textbook moral hazard: analysts take risks in making recommendations, but do not fully internalize the costs of those risks.[2] This could be unproblematic, if quack analysts were quickly driven out of the market. But empirical as well as anecdotal evidence suggests otherwise.[3] This paper formally investigates financial markets' performance at sorting the pundits from the quacks.

It is a well-known fact that markets can learn about analysts' ability based on public information regarding underlying assets.[4] This wisdom is important, but incomplete. Public information varies across time and asset classes. Growth stocks, for instance, may take years before generating revenue. Financial markets must rely then on the private information of investors, in order to learn about analysts' ability.[5] The question is then: Do asset prices aggregate investors' private information about financial analysts? Our paper shows, alas, that this information is (if at all) only properly aggregated in the long run: i.e., when that information matters least.

---

[1] According to R.J. Wayman, Vice President and Portfolio Manager for Sweetwater Asset Management: "Independent research firms are becoming the main source of information on the majority of stocks". `http://www.investopedia.com/articles/analyst/03/031803.asp`.

[2] In short, analysts lack 'skin in the game'. Potential conflicts of interests is, in practice, another important problem. We abstract here from that issue. The interested reader is referred to Ramnath et al. (2008) for an extensive survey of that literature.

[3] The presence of persistent differences in the quality of analyst recommendations has been found in, for instance, Li (2005). Furthermore, pundits in television shows that make stock recommendations are often noted to be poor predictors of stock returns, although they have a significant short-term effect on trading. See for instance `http://en.wikipedia.org/wiki/Wall_Street_Week#The_Rukeyser_Effect`.

[4] See, e.g., Trueman (1994). To the extent that public information is available regarding underlying assets, comparing the predictions of analysts with the realized values of forecasted variables provides an immediate way of evaluating analyst ability.

[5] Investors' private information about analysts is exogenous in the model we consider. Throughout, we remain agnostic about the precise source of this information. Personal contacts is one possibility. The quality of the analyst's reports is another plausible source of information (Hirst et al. (1995) and Asquith et al. (2005) find evidence in this sense). For instance, an investor may discover that an analyst has copied-pasted the arguments he uses, or find flaws in the analyst's underlying reasoning; alternatively, the investor may deem the arguments laid forward by the analyst perceptive and original.

The basic model we analyze has the following features. A finite number of assets of unknown fixed values are traded in a market with a single long-run player (*the analyst*) and an infinite sequence of short-run players (*the investors*). The analyst is of one of two types: informed or uninformed.[6] An informed analyst's recommendations are correlated with the assets' true values. An uninformed analyst, however, knows nothing more about the assets than the public information and makes strategic recommendations aiming to maximize his *reputation*.[7] A new investor arrives each period, possessing two pieces of private information: the latest recommendation from the analyst and an imperfectly informative signal of the analyst's true ability. The investor then decides whether or not to follow the analyst's recommendation, given competitive bid and ask prices. A *liquidity trader* makes his decision based on factors exogenous to the model; a *speculator*, on the other hand, trades to maximize profits. The recommendation and the trade become known publicly when the period ends, at which point the market updates its belief that the analyst is informed (i.e., the reputation of the analyst). If all information could be centralized then the market would learn the true type of the analyst. We explore whether, and how, the private information of investors concerning the analyst is aggregated in the decentralized market setting described above.

Unlike all previous studies, feedback about the analyst's ability is endogenous.[8] Two channels in our model transmit information to the market about the analyst's true type: (i) investors' decision to follow or not the analyst's advice and, (ii) recommendations themselves. How much information each channel conveys is determined endogenously, and varies with time. We say that a *reputational trap* occurs when both channels are mute and no information about the analyst's true type is transmitted to the market in that period.

Our analysis of the basic model revolves around two principal results:

1. We show that in the short run, and in the medium run, the market can get trapped. Speculators follow blindly the recommendations made, and prices fail to aggregate in-

---

[6]We do not model the source of a good analyst's superior information directly. A good analyst may distinguish himself purely by his ability to analyze data, but could also retrieve information from his network of connections within monitored firms.

[7]The basic model assumes that the analyst maximizes his expected reputation one period ahead. We show in a later section that all our results extend to much broader settings, where the analyst maximizes the expected discounted sum of future income, and income in a given period is an increasing function of the analyst's reputation.

[8]See the literature review for previous studies of financial analysts. In our model, feedback about the analyst's ability occurs endogenously through prices and recommendations made (the public information). In particular, assets' true values are never observed. All the paper's main results remain valid if public information about assets' true values is available. This extension is considered in an online-appendix.

vestors' information about the analyst. The uninformed analyst maintains a lasting reputation, and affects prices durably.

2. As time goes by and recommendations accumulate, arbitrage based on the (inferred) ability of the analyst becomes profitable again. Thus, whereas the market gets trapped early on, it is able to sort in the long run the pundits from the quacks.

The basic intuition behind the first point is the following. By raising the expected quality of advice, better reputation increases the average foregone profits from ignoring a recommendation. Hence when reputation is high, so too are foregone profits. Above a threshold reputation speculators must then choose, in equilibrium, to follow *all* advice stemming from the analyst.

The occurrence of reputational traps hinges upon three simple conditions: (a) non-zero measure of liquidity traders, (b) imperfectly informed investors, and (c) good analyst reputation. Without condition (a), prices would be 'too' elastic for a trap to occur: a speculator would then be unable to turn a profit from his information. When the condition holds however, a speculator benefits from prices' relative inelasticity, and the expected profit from following the analyst's advice is increasing in the reputation of the analyst. Next, condition (b) ensures the willingness of speculators to attribute negative ability signals to sheer luck, and more so the better the reputation of the analyst. In view of the previous remarks, condition (c) is the final ingredient leading to a reputational trap: good reputation induces speculators to disregard their information concerning the analyst, and blindly follow the recommendations made.

The occurrence of reputational traps naturally places a question mark over the possibility to learn the analyst's true type based solely on the private information of investors. In our model, however, equilibrium conditions are determined by two layers of uncertainty: uncertainty about the ability of the analyst, and uncertainty about the assets' true values. When the market is trapped, it stops learning about the analyst but continues, through new recommendations, to learn about the assets. This, in turn, transforms the incentives of speculators and allows us to show that reputational traps are transient events.

The intuition is the following. The recommendations of the informed analyst are correlated with the assets' true values. So if the analyst is informed, prices will eventually converge. This implies in turn that the uninformed analyst must either make prices converge too, or lose his reputation.[9] However, as prices converge, speculators tend to gain less from trading based on

---

[9]This mechanism is reminiscent of, e.g., Prendergast and Stole (1996). Effectively, the analyst here makes

the analyst's recommendations, all the while standing to gain more from arbitrage based on their private signal of the analyst: the more information contained in the public history, the more critical the knowledge of the true type of the analyst. An overwhelming dominance of 'buy' recommendations, say, will push prices up toward their highest values. On the other hand, a speculator with a negative ability signal will see his valuations of the assets revert toward the unconditional mean. The higher the prices, the more the trader believes the assets to be overvalued. As prices get close to their upper bound, this will induce him to trade *against* historical trends (i.e. sell the assets), independently of the recommendation made. By acting as contrarians in the long run, speculators thus eventually release information about the true type of the analyst. We go on to show that – in the basic model – the market eventually learns the true type of the analyst.[10]

As illustrated in the narrative above, *learning about analysts is inseparable from the convergence of prices* in the framework we explore. This key and novel feature of our model has important implications. We show for instance in an extension of the basic model that, if asset fundamentals are sufficiently volatile, then learning about analysts will collapse in the short *and in the long run.* The intuition is simple. In the basic model, a trap ends as historical information embedded in prices eventually takes precedence over the current recommendation of the analyst. This in turn forces investors to start turning to their private information with a view to evaluate the credence of this historical information. By contrast, if fundamentals are sufficiently volatile, then current recommendations can remain important relative to historical information. The market may in this case remain trapped infinitely long.[11]

To derive empirical predictions, we then endogenize the decision of the analyst to make a recommendation at all, in any given period. First, we show that in this case the rate of recommendations evolves with time. The intuition is as follows. In general, the uninformed analyst will prefer to avoid making recommendations, so as to minimize the information released concerning his true ability. In a reputational trap however, the analyst knows that the market will follow his recommendations with probability one. He then needs not fear making recommendations. The rate of recommendations thus rises in a reputational trap, and falls outside it. Since reputational traps occur when reputation is good *and* prices moderate,

---

'persistent' recommendations in an attempt to convince the market that he is informed.

[10]The analysis of the long run is complicated in our model by the fact that there is no uniform positive lower bound on the information conveyed about the analyst in any given round: it may be nothing (as in the case of a reputational trap), or close to nothing.

[11]Our long-run learning results thus contrast sharply with those of, e.g., Benabou and Laroque (1992) and Trueman (1994) where learning is entirely unrelated to assets' underlying volatility.

the recommendation rate will increase under these conditions, rather than being monotonic in either of the variables.[12]

Second, our model predicts large and sudden asset price movements. The reason is simple. When the market is trapped, it stops aggregating the information concerning the ability of the analyst, but goes on accumulating information about the assets in the form of new recommendations (at an increased rate moreover, given previous remarks). Large price movements occur when the market starts learning the analyst's true ability again, and suddenly appreciates/depreciates all the information accumulated *during* the reputational trap.

Finally, the price impact from a recommendation is gradual in our model. A recommendation is first incorporated 'at a discount', and later adapted according to the evolution of the analyst's reputation. Past recommendations thus have a contemporaneous effect on prices, so long as the market goes on learning the true ability of the analyst.[13]

Related literature.   This paper examines a previously unexplored aspect of financial markets' (in)efficiency: Do asset prices aggregate investors' private information about analysts? The results in this paper show that strategic – but otherwise uninformed – financial analysts are able to 'manipulate' markets, by exploiting agents' uncertainty regarding their true ability. To this end, we combine elements from two established strands of the economics literature: the market microstructure literature, which examines the ways in which financial markets' working processes determine price formation and trading behavior, and the reputation in repeated games literature, whose concern is the study of continuing relationships (see Vives (2012) and Mailath and Samuelson (2013) for recent and detailed surveys of these vast and rich literatures). We add to the first strand by modelling explicitly the provision of information by a strategic financial analyst. We contribute to the second strand by studying a setting in which exogenous private signals provide the short-run players of our model with noisy information concerning the realized type of the long-run player. Recent contributions by Wiseman (2009) and Hu (2014) explore related but crucially distinct settings, in which exogenous *public* signals are observed regarding the true type of the long-run player.[14]

Our paper is related to the literature on social learning, whereby Bayes-rational individuals learn sequentially from the discrete actions of others; early contributions include Banerjee

---

[12]To the best of our knowledge, no other paper makes predictions about the *rate* of recommendations of financial analysts.

[13]See, e.g., Michaely and Womack (2005) for empirical evidence of this phenomenon.

[14]See also Section 4 for a detailed discussion of our results in relation to the seminal papers of Benabou and Laroque (1992) and Cripps et al. (2004).

(1992), Bikhchandani et al. (1992), and Smith and Sørensen (2000). However, two crucial differences distinguish our paper from the aforementioned. One, information about assets is mediated in our model by another strategic player (the analyst). Two, there is no herding in our model. Investors ignore their private information not because they herd on the actions of others, but because the recommendations of the analyst provide a window of opportunity to make profits. The multidimensional nature of uncertainty is here key, relating our paper to Avery and Zemsky (1998) and Park and Sabourian (2011). In our model, the first layer of uncertainty is related to asset fundamentals. The second layer, on the other hand, pertains to the *quality* of information accumulated concerning the true values of assets. This, and the fact that information is provided by an analyst, distinguishes our paper from those above.

Whereas the amount of empirical work on financial analysts is enormous, theoretical contributions are relatively few.[15] Notable exceptions include Admati and Pfleiderer (1986), Benabou and Laroque (1992), Trueman (1994), and Ottaviani and Sørensen (2006). Admati and Pfleiderer (1986) examine a monopolistic analyst who, in order to overcome the dilution in the value of information due to its leakage through informative prices, may prefer to sell noisier versions of the information he actually has. There is no uncertainty about the ability of the analyst in that model. Similarly, the paper by Benabou and Laroque (1992) studies an altogether different problem than ours. All analysts are informed, but engage in insider trading and have therefore strong short-run incentives to deceive the market. In Trueman (1994) and Ottaviani and Sørensen (2006), good and bad analysts coexist, but trade plays no role (there are no traders). Moreover, public information about underlying assets is plentiful: the asset's true value is observed at the end of a period.[16]

The paper is organized as follows. Section 2 introduces the basic model. Section 3 analyses the short run, and Section 4 the medium and long run. Section 5 extends the model to analyze the role played by the volatility of assets, to derive empirical implications, and to demonstrate the robustness of our results to more general specifications. Proofs not included in the text are presented in an appendix. We discuss further extensions of the model in an online-appendix.

---

[15]The paper by Spiegler (2006) studies, in a different context, a market for quacks. There is no uncertainty about the ability of the expert, and agents are boundedly rational.

[16]The related literature on 'reputational herding' (Scharfstein and Stein, 1990; Dasgupta and Prat, 2008) investigates a separate issue, namely learning about the ability of fund managers (viz. investors).

2 Basic model

The basic model has the following broad features. Time is discrete. There is one long-run player (the analyst) and an infinite sequence of short-run players (the investors).[17] Each period, a new investor is given the option to trade an asset of unknown value based on a recommendation from the analyst. The analyst may be a pundit or a quack. A pundit's recommendations are correlated with the asset's true value, but a quack is uninformed and makes strategic recommendations with a view to maximize his reputation. Each investor possesses a piece of information concerning the true type of the analyst. We now lay out the details and notation of this model.

Assets and Financial Analyst. Let $\mathcal{A}$ denote a finite set of assets, with i.i.d. (fixed) fundamental values $\{\theta_i\}_{i \in \mathcal{A}}$. The fundamental values have mean zero and realizations in $\{-1, 1\}$.[18] Each period t nature publicly selects, uniformly at random, one asset in $\mathcal{A}$ to be traded during that period.[19] Let $\iota(t)$ denote the asset selected. Whenever it is unlikely to create confusion, we shorten notation by using θ_t instead of $\theta_{\iota(t)}$ to indicate the fundamental value of the asset being traded during period t.

Each period t, the analyst makes a recommendation concerning $\iota(t)$. We let $r_t \in \{-1, 1\}$ denote the recommendation made in period t. The analyst is either informed or uninformed ('good' or 'bad'); τ denotes the analyst's true type: $\tau = G$ if he is good, $\tau = B$ if he is bad. His type is drawn at the beginning of the game, fixed henceforth, and known to himself only: $\tau = G$ with probability $\lambda_0 \in (0, 1)$, and $\tau = B$ with complementary probability $1 - \lambda_0$. The parameter λ_0 defines the *reputation* of the analyst at the beginning of the game.

The informed analyst makes recommendations correlated with the true asset values:[20]

$$\mathbb{P}(r_t = \theta_t | \tau = G) = \phi,$$

---

[17]The market microstructure is a dealer model adapted from Glosten and Milgrom (1985).

[18]As in Glosten and Milgrom (1985), one may think of θ_i as the expected sum of discounted dividend payments, where the dividend d_{it} in period t takes a value in $\{-1, 0, 1\}$. The frequency φ_i of dividend payments, and the relative frequency of high/low dividends then together determine θ_i. As $\varphi_i \to 0$, public information about θ_i vanishes. See also the online appendix for a model with public information about assets' fundamental values.

[19]This assumption is made for simplicity. Alternative assumptions complicate the analysis but leave the main results unchanged. For instance, the investor in period t could choose $\iota(t)$, or the analyst.

[20]In the terminology of the reputation literature, the informed analyst is a behavior type. Implicit in our formulation is the assumption that the good analyst receives a sequence of signals correlated with the assets' true values, and commits to reveal those signals truthfully. In Rudiger and Vigier (2014) we allow the good analyst to be strategic as well, and show that there is an equilibrium in which he is truthful.

$\phi \in (1/2, 1)$. The recommendations of the good analyst are conditionally independent across time periods.

The bad analyst knows the information publicly available, but nothing more than that, and makes strategic recommendations with a view to maximize his reputation next period.[21] The analyst's reputation is updated using all public information (we later describe the updating process).[22]

Investors. A new investor arrives to the market each period.[23] The investor knows all public information (later specified in details), and two more pieces of information: the recommendation from the analyst, r_t, and a signal $s_t \in \{0, 1\}$ of the analyst's true ability distributed according to[24]

$$\mathbb{P}(s_t = 0 | \tau = G) = 1 - \pi,$$
$$\mathbb{P}(s_t = 0 | \tau = B) = 1.$$

These signals are conditionally independent across investors. The parameter π captures the precision of the signal s_t. If π were zero then investors would always follow the recommendations made. We rule out this trivial and uninteresting scenario, by assuming $\pi > 0$. Observe too that if $\pi = 1$ then each investor knows the true type of the analyst.[25]

An investor is of one of two types, defining his motives for trade. With probability $\mu \in (0, 1)$ investor t is a *speculator*, maximizing expected profits from trade. With remaining probability $1 - \mu$ investor t is a *liquidity trader*, trading for exogenous motives unrelated to profits. Investor types are independent across time periods. A liquidity trader buys, abstains or sells the asset with probability $1/3$ each. We let $y_t \in \{a, n, b\}$ denote the trade of investor t: $y_t = a$ if he buys the asset, $y_t = n$ if he abstains, and $y_t = b$ if he chooses to sell the asset.[26]

---

[21]This assumption is relaxed in Section 5.3 where we show that our results hold, qualitatively, when the analyst maximizes the expected discounted sum of future income, and income in a given period is an increasing function of the analyst's reputation.

[22]Chen et al. (2005) find that the market's response to analysts' recommendations is consistent with investors learning about analysts' forecasting ability in a Bayesian fashion as more observations of past recommendations become available.

[23]As in Glosten and Milgrom (1985) investors are short-lived, for tractability.

[24]More general signal structures are considered in the online appendix.

[25]As indicated earlier, we remain largely agnostic in this paper about the source of the information summarized in s_t. At any rate, we make no assumption about the strength of the signal s_t, other than $\pi > 0$.

[26]As in Glosten and Milgrom (1985) the size of trades is irrelevant, since speculators are risk neutral. It is assumed that each trade involves one unit of the asset.
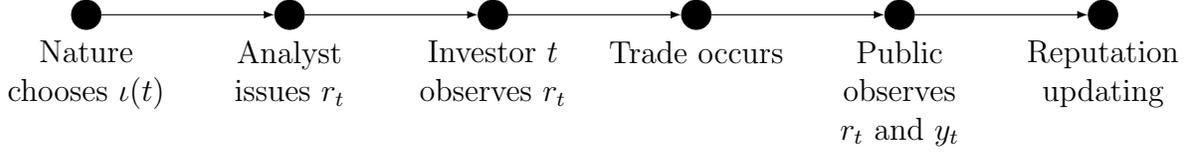
Figure 1: Timing within period t

Timing. The timing of the game within each period is as follows (c.f. Figure 1). The asset $\iota(t)$ is first determined. The analyst then makes his recommendation, r_t. Investor t observes r_t, and chooses whether or not to follow this recommendation given the bid and ask prices posted.[27] This investor then leaves the market. The recommendation r_t and trade y_t become known publicly when the period ends. This timing assumption is crucial, but reflects a real-life feature: analyst recommendations are first disclosed to client investors before being publicized (see, e.g., Michaely and Womack (2005)).[28] The market then revises its beliefs, and updates accordingly the prices posted for the next period.

Some notation is useful. We let $\mathcal{H}_t$ denote the public information at the beginning of period t:

$$\mathcal{H}_t := \{\mathcal{R}_{t-1}, \mathcal{Y}_{t-1}, \mathcal{I}_t\},$$

where $\mathcal{R}_{t-1} := (r_0, ..., r_{t-1})$, $\mathcal{Y}_{t-1} := (y_0, ..., y_{t-1})$, and $\mathcal{I}_t := (\iota_0, ..., \iota_t)$.

We let $\mathbb{P}_t(\cdot) := \mathbb{P}(\cdot | \mathcal{H}_t)$ denote the probability operator conditional on $\mathcal{H}_t$, and $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | \mathcal{H}_t]$. The reputation of the analyst at time t, λ_t, is defined as the probability that the analyst is informed, given all public information at the beginning of period t:

$$\lambda_t := \mathbb{P}_t(\tau = G).$$

The (Bayesian) reputation updating equation is

$$\lambda_{t+1}(r_t, y_t) := \frac{\lambda_t \mathbb{P}_t(r_t, y_t | G)}{\lambda_t \mathbb{P}_t(r_t, y_t | G) + (1 - \lambda_t)\mathbb{P}_t(r_t, y_t | B)}. \tag{1}$$

---

[27]To avoid unnecessary cluttering, we do not model market makers explicitly in this paper. Bid and ask prices are specified in (2).

[28]In fact, powerful investors typically become aware of analysts' recommendations before these are even disclosed to clients. See, e.g., `http://www.nytimes.com/2014/01/09/business/blackrock-agrees-to-stop-pursuing-nonpublic-views.html`, or Rudiger and Vigier (2015) for a related theoretical study.

9

Strategies. The strategic players of our model are (i) the analyst of type $\tau = B$, and (ii) the speculators. The vector $\underline{\sigma}_{\mathcal{H}_t}$ denotes the strategy of the uninformed analyst in period t given history $\mathcal{H}_t$, with $\sigma^r_{\mathcal{H}_t}$ indicating the probability of recommending r. The vector $\underline{\xi}_{\mathcal{H}_t}$ denotes a speculator's strategy in period t given $\mathcal{H}_t$, with $\xi^y_{\mathcal{H}_t}(r_t, s_t)$ denoting the probability that he takes action y having observed recommendation r_t and signal s_t. To shorten notation, and when this is unlikely to create confusion, we will use $\underline{\sigma}_t, \underline{\xi}_t$ instead of $\underline{\sigma}_{\mathcal{H}_t}, \underline{\xi}_{\mathcal{H}_t}$.

Equilibrium. Prices are assumed competitive, reflecting all information available publicly at the time of trade. This includes the trade order itself, and the behavior of all strategic players:

$$p^y_t = \mathbb{E}_t[\theta_t | y_t = y; \underline{\xi}_t, \underline{\sigma}_t]. \tag{2}$$

The structure of the game described above is common knowledge. A *competitive market equilibrium* is defined by sequences $\left(\underline{\xi}_{\mathcal{H}_t}\right)_{\mathcal{H}_t}$ and $\left(\underline{\sigma}_{\mathcal{H}_t}\right)_{\mathcal{H}_t}$ maximizing expected profits for speculators and expected reputation for the analyst of type $\tau = B$, with prices given by (2).

Definition 1. *The sequences $\left(\underline{\xi}_{\mathcal{H}_t}\right)_{\mathcal{H}_t}$ and $\left(\underline{\sigma}_{\mathcal{H}_t}\right)_{\mathcal{H}_t}$ constitute a **competitive market equilibrium** if and only if*

- $\xi^a_t(r_t, s_t) > 0 \Rightarrow \mathbb{E}_t[\theta_t | r_t, s_t; \underline{\sigma}_t] \geq p^a_t$, *with p^a_t given by (2).*

- $\xi^b_t(r_t, s_t) > 0 \Rightarrow \mathbb{E}_t[\theta_t | r_t, s_t; \underline{\sigma}_t] \leq p^b_t$, *with p^b_t given by (2).*

- $\mathbb{E}_t[\theta_t | r_t, s_t; \underline{\sigma}_t] > p^a_t \Rightarrow \xi^a_t(r_t, s_t) = 1$, *with p^a_t given by (2).*

- $\mathbb{E}_t[\theta_t | r_t, s_t; \underline{\sigma}_t] < p^b_t \Rightarrow \xi^b_t(r_t, s_t) = 1$, *with p^b_t given by (2).*

- $\sigma^r_t > 0 \Rightarrow r \in \arg\max_{r_t} \mathbb{E}_t[\lambda_{t+1}(r_t, y_t) | \tau = B; \underline{\sigma}_t, \underline{\xi}_t]$, *with $\lambda_{t+1}(r_t, y_t)$ given by (1).*

Part one (resp. part two) says that speculators do not buy the asset (resp. sell the asset) if they expect to lose money by doing so. Part three (resp. part four) says that speculators always buy the asset (resp. sell the asset) if by doing so they expect to make strictly positive profits. Part five says that an analyst of type $\tau = B$ only issues a recommendation if no other recommendation (on average) yields higher reputation next period.

3   The short run

We explore in this section the basic workings of our model and show that the market gets trapped in the short run when the analyst is reputable, failing to properly aggregate investors' private information concerning the former's ability.

Let θ_t^G denote asset $\iota(t)$'s expected value in period t, assuming the analyst is informed. Let $\theta_t^G(r)$ denote the corresponding expected value, where we augment the information set with the recommendation made in period t. Thus

$$\theta_t^G := \mathbb{E}_t[\theta_t | \tau = G],$$
$$\theta_t^G(r) := \mathbb{E}_t[\theta_t | r_t = r, \tau = G].$$

Let q_t^r be the probability that the informed analyst recommends r in period t,

$$q_t^r := \mathbb{P}_t(r_t = r | \tau = G).$$

Straightforward calculations yield[29]

$$q_t^+ = \frac{1}{4}\left[1 + (2\phi - 1)\theta_t^G\right]$$
$$q_t^- = \frac{1}{4}\left[1 - (2\phi - 1)\theta_t^G\right].$$

Consider next the speculators. An investor possesses two pieces of information with which to make profits: the latest recommendation, and his private signal of the analyst's ability. Let $v_t(r, s)$ denote the valuation of an investor with recommendation r and ability signal s. An investor with $s_t = 1$ assigns probability 1 to the informed analyst. This yields

$$v_t(1, 1) = \theta_t^G(1),$$
$$v_t(-1, 1) = \theta_t^G(-1).$$

An investor with $s_t = 0$, on the other hand, updates his assessment of the analyst in two ways: first according to the likelihood $(1 - \pi)/1$ of a negative ability signal, and second according

---

[29]We use the short-hand notation q_t^+ and q_t^- for q_t^{+1} and q_t^{-1}. Similarly for σ_t^r.

to the likelihood $q_t^{r_t}/\sigma_t^{r_t}$ of the recommendation made; this yields

$$v_t(1,0) = \frac{\lambda_t q_t^+(1-\pi)}{\lambda_t q_t^+(1-\pi) + (1-\lambda_t)\sigma_t^+}\theta_t^G(1),$$

$$v_t(-1,0) = \frac{\lambda_t q_t^-(1-\pi)}{\lambda_t q_t^-(1-\pi) + (1-\lambda_t)\sigma_t^-}\theta_t^G(-1).$$

We next examine the ranking of the valuations derived above. Much of our model's interest springs from the fact that rather than being fixed, the ranking of the valuations – and, by way of consequence, the information reflected in prices – typically evolves over time. To help fix ideas we focus in what follows on histories where $\theta_t^G(-1) > 0$, which we will refer to as cases of 'bullish' history.[30]

Claim 4 in the appendix establishes that, if history is bullish, then in any equilibrium:

$$v_t(-1,0) < \min\{v_t(-1,1), v_t(1,0)\}, \tag{3}$$

$$v_t(1,1) > \max\{v_t(-1,1), v_t(1,0)\}. \tag{4}$$

The intuition is straightforward. An investor with $(r_t = -1, s_t = 0)$ discounts the bullish history *and* receives a recommendation to sell the asset. At the other extreme, an investor with $(r_t = 1, s_t = 1)$ endorses the bullish history *and* receives a recommendation to buy the asset. Following (3)-(4), these investors respectively have the lowest and highest valuations.

We deduce from the former inequalities that, in equilibrium, a speculator with $(r_t = 1, s_t = 1)$ (resp. $(r_t = -1, s_t = 0)$) must buy the asset (resp. sell the asset) with probability 1. Speculators with $(r_t = 1, s_t = 0)$ or $(r_t = -1, s_t = 1)$, on the other hand, face a dilemma. A speculator with $(r_t = 1, s_t = 0)$, for instance, discounts the bullish history, but is simultaneously advised to buy the asset. In general, his trading decision will thus vary according to the relative weight attached to these two factors. The implications of his decision, however, are crucial for the transmission of information concerning the analyst. If he abstains from buying, the market will be able to distinguish him (statistically) from the investor with $(r_t = 1, s_t = 1)$, and thereby learn something about s_t and hence the analyst's true type. If he always buys the asset, however, the market will learn nothing concerning the analyst.

Henceforth, we will say that the market is *trapped* at time t if, in the equilibrium considered and given the history, the market fails to aggregate any information about the true type of

---

[30]This situation will occur anytime the total number of 'buy' recommendations on asset $\iota(t)$ is at least two more than the total number of 'sell' recommendations on that asset.

the analyst during that period.

Definition 2. *The market is **trapped** at time t if*

$$\mathbb{P}_t(\lambda_{t+1} = \lambda_t) = 1. \tag{5}$$

Extending the previous definition to focus first on the short run, we will say that the market is trapped for T initial periods if in the equilibrium considered $\mathbb{P}(\lambda_T = \lambda_0) = 1$. Our first theorem shows that the market will be trapped for an arbitrarily long time provided the initial reputation of the analyst is good enough.

The basic intuition is the following. Holding fixed the strategy of speculators, better reputation increases (i) the expected quality of advice and (ii) the amount of information revealed by trades. We show that the first effect dominates, such that – in equilibrium – better reputation increases average foregone profits from ignoring a recommendation. Hence when reputation is high, so too are foregone profits. Above a threshold reputation speculators must then choose, in equilibrium, to follow *all* advice stemming from the analyst.

Theorem 1. *If $\mu < 1$ and $\pi < 1$ then there exists a strictly increasing sequence $\{\hat{\lambda}_0^T\}_{T \in \mathbb{N}}$ with limit 1, such that the following property holds: the market is trapped, in any equilibrium, for T initial periods if and only if $\lambda_0 \geq \hat{\lambda}_0^T$. The greater the mass of speculators (μ) and/or the greater the precision of speculators' information (π) the higher $\hat{\lambda}_0^T$, for all T.*

The occurrence of a trap hinges upon three simple conditions: (a) non-zero mass of liquidity traders ($\mu < 1$), (b) imperfectly informed investors ($\pi < 1$), and (c) good analyst reputation ($\lambda_0 \geq \hat{\lambda}_0^T$). Condition (a) bounds the elasticity of prices. Consider a speculator with a positive recommendation, looking to buy the asset. In the absence of liquidity traders, a 'lemons' situation ensues in which the market unravels: prices adjust, and only the speculator with the highest valuation would ever buy the asset. But a speculator's trade order would then reveal his information of the analyst since, as we saw earlier, a speculator with ($r_t = 1, s_t = 1$) always has the highest valuation. Condition (a) is thus necessary for the market to get trapped.[31] Conditions (b) and (c) are complementary. When they hold, speculators willingly attribute negative ability signals to sheer luck. This allows them to disregard their information concerning the analyst, and blindly follow the recommendations made.[32]

---

[31]When $\mu = 1$ there is also a 'no trade' equilibrium in which the market breaks down, and a trap trivially ensues. Condition (a) is necessary for the result to hold in *all* equilibria.

[32]The break-down of learning about the expert's ability occurring in our model is related to the market

The results on comparative statics are intuitive. The less the mass of liquidity traders, the more responsive the prices to the trade orders. This in turn discourages speculators from ignoring their signals of the analyst's ability, preventing thereby the occurrence of reputational traps. Similarly, an increase in the parameter π raises the precision of investors' information regarding the analyst, thus discouraging the former from ignoring that information.

The remainder of this section elaborates the various steps leading to Theorem 1, and derives closed-form expressions of the thresholds $\{\hat{\lambda}_0^T\}_{T \in \mathbb{N}}$ defined in the statement of the theorem. We first construct an equilibrium in which, when the reputation of the analyst is good, the market fails to aggregate information about τ from one period to the next. The second step is to show that the threshold reputation found in step one is in fact 'uniform' across all equilibria, i.e. when reputation is above the threshold then the market *must* fail to learn about τ, in all equilibria. This step is hard; the reader who wishes to skip the details of the analysis can go directly to the next section.

Two channels in our model transmit information in equilibrium to the market about the analyst's true type. First, speculators' decision to follow or not the analyst's advice. Second, recommendations themselves. Information about the analyst's ability is transmitted through the recommendations because the extent to which speculators follow the analyst's advice will in general not be uniform across the different possible recommendations. Therefore, as the bad analyst wishes to minimize the information channeled regarding his type, he will favor the recommendations most likely to be followed by the speculators. Thus, in equilibrium the recommendations themselves will generally be informative about the true type of the analyst.

The next definitions help formalize the distinction made between the two channels identified above.

Definition 3. *Given trading strategy $\underline{\xi}_t$, say that **screening is efficient** if trades fully reveal private signals of the analyst's ability, in the following sense: $\xi_t^y(r, 1) > 0 \Rightarrow \xi_t^y(r, 0) = 0$, $\forall\, y, r$.*

---

break-down occurring in Ely and Välimäki (2003). In their model a series of Principals sequentially interact with a single Agent, whose type (good or bad) is unknown to the Principals. In some cases, the good Agent should take a given action, but that action is also the preferred action of a bad Agent. To distinguish himself, the good Agent then favors the inefficient action, but by doing so he kills the Principals' incentive to hire him. In both that paper and ours, the crux lies in the failure of each Principal (viz. each speculator) to internalize the benefits to others from learning about the type of the Agent (viz. the analyst). However the break-down of learning takes opposite forms in the two papers: In Ely and Välimäki (2003) learning breaks down because the Principals stop 'trading'; by contrast learning breaks down (viz. a reputational trap occurs) in our paper because the Principals trade with probability one.

Definition 4. *Given trading strategy $\underline{\xi}_t$, say that **screening breaks down** if trades are uninformative about private signals of the analyst's ability, in the following sense: screening breaks down on the positive side (resp., negative side) if there exist $y \in \{a, n, b\}$ such that $\xi_t^y(1,1) = 1 = \xi_t^y(1,0)$ (resp., such that $\xi_t^y(-1,1) = 1 = \xi_t^y(-1,0)$). Say that screening breaks down if it does so on both the positive and the negative side.*

Suppose next an equilibrium and history $\mathcal{H}_t$ exist such that the market is trapped at time t. In this equilibrium and given that history, the market fails to aggregate information about τ during period t. Observe then that screening must break down during that period. If it did not, then the investor's trade in period t would act as a (noisy) signal of information s_t, and (5) would not hold. The next observation establishes the converse of this result.

Observation 1. *Consider an equilibrium and history $\mathcal{H}_t$ where screening breaks down in period t. Then (5) holds in this equilibrium, given that history.*

To see why the observation is true, suppose for the sake of contradiction that (5) does not hold. The recommendation made in period t must then be informative about the analyst, i.e. the equilibrium must be such that $\underline{\sigma}_t \neq \underline{q}_t$. Let r and r' such that $\sigma_t^r > q_t^r$ and $\sigma_t^{r'} < q_t^{r'}$. In that case (1) yields $\lambda_{t+1}(r, y) < \lambda_t$ and $\lambda_{t+1}(r', y) > \lambda_t$, for all y. Recommending r' is thus unambiguously more advantageous for the uninformed analyst. This is inconsistent with the uninformed analyst maximizing his expected reputation, and so (5) must hold, by contradiction.

Observation 1 stresses the strategic implications of our model. Consider for comparison the case of an uninformed analyst mechanically (uniformly) randomizing between buy and sell recommendations, i.e. behaving according to $(\sigma_t^-, \sigma_t^+) = (1/2, 1/2)$ regardless of the history. Suppose the history were bullish: the market would then expect the informed analyst to be bullish too. But then, an analyst who emits buy and sell recommendations with equal probability will (on average) reveal himself by appearing bearish, comparatively. Thus, in that case, in contrast to our model, screening may break down and yet some information be conveyed in equilibrium to the market about the true type of the analyst.

We now ask: When would speculators prefer ignoring their information concerning the

analyst, in a competitive market equilibrium? The following notation will be used:

$$\hat{\lambda}^-(\mathcal{H}_t) := 1 - \frac{(1-\gamma)\left(1 - \frac{\theta_t^G(-1)}{\theta_t^G}\right)}{\gamma q_t^{-\frac{\theta_t^G(-1)}{\theta_t^G}} + 1 - \gamma}, \tag{6}$$

$$\hat{\lambda}^+(\mathcal{H}_t) := 1 - \frac{1-\pi}{\pi} \cdot \frac{(1-\gamma)\left(1 - \frac{\theta_t^G}{\theta_t^G(1)}\right)}{\gamma q_t^+ + (1-\gamma)\frac{\theta_t^G}{\theta_t^G(1)}}, \tag{7}$$

and

$$\hat{\lambda}(\mathcal{H}_t) = \max\{\hat{\lambda}^+(\mathcal{H}_t), \hat{\lambda}^-(\mathcal{H}_t)\}.$$

Notice that the right-hand side of these expressions are functions of $\mathcal{H}_t$ through the θ_t^G's.[33]

Lemma 1 shows that if in a given period reputation is above the threshold $\hat{\lambda}(\mathcal{H}_t)$ then an equilibrium can be found in which the market learns nothing about τ during that period. The proof is simple, and helps illustrate the underlying forces at work; we therefore keep it in the text.

Lemma 1. *If in some equilibrium, given history $\mathcal{H}_t$, $\lambda_t \geq \hat{\lambda}(\mathcal{H}_t)$ then an equilibrium exists in which (5) holds, given $\mathcal{H}_t$. Moreover, $\hat{\lambda}(\mathcal{H}_t)$ is the smallest number satisfying this property.*

Proof of Lemma 1: Continue assuming history is bullish. We noted earlier that when history is bullish then, in equilibrium, a speculator with $(r_t = -1, s_t = 0)$ sells the asset with probability 1. Similarly, a speculator with $(r_t = 1, s_t = 1)$ buys the asset with probability 1. Screening thus breaks down if and only if (a) speculators with $(r_t, s_t) = (1, 0)$ buy with probability 1, while (b) speculators with $(r_t, s_t) = (-1, 1)$ sell with probability 1. We next derive the necessary and sufficient conditions for (a) and (b) to hold in equilibrium.

Given competitive prices – and in view of Observation 1 – (a) is an equilibrium if

$$\frac{\lambda_t(1-\pi)q_t^+}{\lambda_t(1-\pi)q_t^+ + (1-\lambda_t)q_t^+} \cdot \theta_t^G(1) \geq \frac{\gamma q_t^+}{\gamma q_t^+ + 1 - \gamma} \cdot \lambda_t \theta_t^G(1) + \frac{1-\gamma}{\gamma q_t^+ + 1 - \gamma} \cdot p_t^0,$$

where $\gamma := \frac{\mu}{\mu + (1-\mu)/3}$ and $p_t^0 := \mathbb{E}_t[\theta_t]$. The left-hand side of this inequality is $v_t(1, 0)$, when $\sigma_t^+ = q_t^+$. The right-hand side is obtained as follows. With probability μq_t^+ investor t is a speculator observing $r_t = 1$. With probability $(1-\mu)/3$, he is a liquidity trader looking to

---

[33]Note too that, given a bullish history, $\hat{\lambda}^-(\mathcal{H}_t) \in (0, 1)$ and is, moreover, independent of π. By contrast $\hat{\lambda}^+(\mathcal{H}_t)$ is strictly monotonic in π, tending to 1 as π tends to 1, and to $-\infty$ as π tends to 0. Thus we can find $\hat{\pi}(\mathcal{H}_t)$ such that $\hat{\lambda}(\mathcal{H}_t) = \hat{\lambda}^-(\mathcal{H}_t)$ if $\pi \leq \hat{\pi}(\mathcal{H}_t)$, and $\hat{\lambda}(\mathcal{H}_t) = \hat{\lambda}^+(\mathcal{H}_t)$ otherwise.

buy the asset. In the former case the expected asset value is $\lambda_t \theta_t^G(1)$. In the latter case, it is simply the public valuation p_t^0. We find, by substituting $p_t^0 = \lambda_t \theta_t^G$, that the inequality above is satisfied for $\lambda_t \geq \hat{\lambda}^+(\mathcal{H}_t)$.

Similarly, for (b) to be an equilibrium requires

$$\theta_t^G(-1) \leq \frac{\gamma q_t^-}{\gamma q_t^- + 1 - \gamma} \cdot \lambda_t \theta_t^G(-1) + \frac{1 - \gamma}{\gamma q_t^- + 1 - \gamma} \cdot p_t^0.$$

Substituting again $p_t^0 = \lambda_t \theta_t^G$, this inequality is satisfied for $\lambda_t \geq \hat{\lambda}^-(\mathcal{H}_t)$. The proof is similar for other histories.

∎

Our next result shows that $\hat{\lambda}(\mathcal{H}_t)$ is in fact a uniform bound, in the following sense:

Lemma 2. *Consider an equilibrium and history $\mathcal{H}_t$ such that $\lambda_t \geq \hat{\lambda}(\mathcal{H}_t)$. Then (5) holds in this equilibrium, given that history.*

The proof of the lemma contains various steps. We here sketch the main arguments.[34] Let $E1$ denote the trapped equilibrium of Lemma 1, and suppose another equilibrium exists, $E2$ say. We proceed to analyze the second equilibrium.

We omit in this discussion the time subscripts, to avoid cluttering notation. Denote by $\underline{\sigma}_1$ and $\underline{\sigma}_2$ the strategy played by the uninformed analyst in these equilibria, respectively. First, suppose that $\sigma_2^+ > \sigma_1^+$. The expected reputation from issuing a positive recommendation must then be lower in $E2$ than in $E1$: one, a positive recommendation is more likely to be sent by the bad analyst in $E2$ than in $E1$; two, there can be no less screening in $E2$ than in $E1$, since in $E1$ screening altogether breaks down.

The key step of the proof is then as follows. We first show that *the less often a recommendation is issued by the uninformed analyst, the less speculators gain from screening it.*[35] Thus: since no screening occurs in $E1$, and $\sigma_2^- < \sigma_1^-$, screening must break down on the negative side in $E2$. But then the expected reputation from issuing a sell recommendation is greater in $E2$ than it is in $E1$. Combining this observation with our first remark above finally shows that in $E2$, the expected reputation from recommending $r_t = -1$ must be greater than the expected reputation from recommending $r_t = 1$. But this is inconsistent with equilibrium.

---

[34]A detailed proof is given in the appendix.

[35]The statement is intuitive and indeed, if prices are fixed, proving it is trivial. The difficulty comes from the fact that prices are competitive in our model.

These arguments thus establish $\underline{\sigma}_1 = \underline{\sigma}_2$. The proof of Lemma 2 is concluded by showing that in $E2$, screening breaks down too, and that prices in the two equilibria are identical. The two equilibria considered thus coincide in that period, given that history.

A proof of Theorem 1 is obtained by combining Lemmas 1 and 2, and exploiting certain regularity properties of the thresholds $\hat{\lambda}(\mathcal{H}_t)$ with respect to the history of recommendations. The details are in the appendix.

4   The medium and long run

We examine in this section the medium and long-run properties of our model. We first establish that the reputational traps uncovered in Section 3 are transient events. We then show that even if it exits a trap, the market may be trapped again at a later date. In spite of this, the section's main result shows that the market learns the ability of the analyst in the long run.

The occurrence of reputational traps naturally places a question mark over the possibility to learn τ and, in turn, over prices' convergence to the fundamental asset values. However, we next argue that whenever the market is trapped, a countervailing force starts working. Recall that there are two layers of uncertainty in the market we model: uncertainty about the ability of the analyst, τ, and uncertainty about the asset values, $\{\theta_i\}_{i \in \mathcal{A}}$. When the market is trapped, it stops learning τ, but goes on learning $\{\theta_i\}_{i \in \mathcal{A}}$. This transforms gradually the incentives of speculators. As a result:[36]

Proposition 1. *Reputational traps are transient events, almost surely, in any competitive market equilibrium.*

The basic intuition behind Proposition 1 is as follows. When the market is trapped, the uninformed analyst can hide his type insofar as he pretends to be good. If he were good, however, then prices would converge to the assets' true values. This forces the analyst of type $\tau = B$ to make prices converge too, or else lose his reputation.[37] However, as prices converge, speculators tend to gain less from trading based on the recommendations and, all the while, stand to gain more from arbitrage based on their information of the analyst's true ability. To be sure, the signal s_t determines the 'discount factor' of *all* past recommendations. Hence, the more information contained in the public history, the more critical the signal of the analyst's

---

[36]The formal arguments for this proposition are contained in the proof of Theorem 2.

[37]Effectively, the uninformed analyst makes his recommendations persistent in an attempt to pool with the informed analyst. A related mechanism is studied in Prendergast and Stole (1996).

true type. An overwhelming dominance of 'buy' recommendations, say, will push prices up toward their highest value. A speculator with a negative ability signal on the other hand will see his valuations of the assets revert toward the unconditional mean (i.e zero). Hence, the higher the prices the more the speculator will view the assets as overvalued. As prices approach their upper bound, this will induce him to trade *against* historical trends (i.e. sell the assets), independently of the recommendations made. By acting as contrarians (relative to the history), speculators thus eventually release information about the true type of the analyst. This establishes that the reputational traps uncovered in Section 3 are transient events.

Perhaps unfortunately, we now show that anytime one trap ends, the market will with positive probability re-enter another one.[38]

Proposition 2. *Assume $\mathcal{A}$ has two elements at least. In any equilibrium, given any history, the probability that the market will be trapped some time in the future is positive (strictly).*

The arguments establishing Proposition 2 reverse those used for Proposition 1. If the analyst becomes sufficiently reputable, then the information incorporated in his current recommendation will swamp that contained in a speculator's signal of the analyst's ability. Each period, a threshold exists such that if reputation is above that threshold, the market will be trapped (Lemmas 1 and 2). Furthermore, this threshold depends on past recommendations made *on that period's asset* only. Random movements in reputation related to trades in other assets ensure that it is always possible that reputation will move above this threshold.

We go on to show that in the basic model – and in spite of the last result – the market learns the true ability of the analyst in the long run.[39] The key to our second theorem is to show that almost surely screening becomes *efficient* eventually.[40]

Theorem 2. *The market almost surely learns the analyst's true type in the long run, in any equilibrium. In particular:*

1. *Conditional on type $\tau = G$, prices converge almost surely to the true asset values.*

2. *Conditional on type $\tau = B$, prices converge almost surely to zero.*

---

[38]The case where the market has a single asset is unresolved. We have neither been able to prove nor disprove the result in that case.

[39]In the model we explore, small frictions due to transaction costs would not affect learning, since analyst recommendations here accumulate whether or not trade occurs. This observation remains true even when the strategic analyst is given the option not to make recommendations (see Section 5.2).

[40]This intermediary step is shown in Claim 9 of the appendix.

Our analysis of the long run provides interesting complementary insights relative to the seminal work of Benabou and Laroque (1992) and Cripps et al. (2004). In both papers, a long-run player is faced with a sequence of short-run players. The long-run player is either a commitment type – committed to playing action A, say – or a strategic type, with a choice of playing action D. The long-run player's action is unobserved, but generates a public signal y distributed according to F^A if he takes action A and F^D if he takes action D.[41] The short-run players' best response to A is α, say, but (A, α) is not a Nash equilibrium of the stage-game. As $F^D \neq F^A$, the strategic type thus eventually reveals himself by deviating to play action D, when his reputation is good enough. By contrast, in our paper (A, α) *is* a Nash equilibrium of the stage-game, when the reputation of the analyst is good enough. In those cases, it is prices' convergence which ultimately ensures that (A, α) *does not remain* a Nash equilibrium.[42] Learning about analysts is thus inseparable from the convergence of asset prices. This key and novel feature of our model has important implications. It underlies, for instance, Proposition 3 of the next section, where we show that if asset fundamentals are sufficiently volatile then learning about analysts will collapse in the short *and in the long run*.

5   Some extensions of the basic model

We extend in this section the basic model along three broad directions, with a view to explore some issues related to the central theme of our paper. First, Section 5.1 examines the impact of volatile fundamentals. Section 5.2 endogenizes the decision to make new recommendations. Section 5.3, finally, investigates the effects from expanding the time horizon of the uninformed analyst. Other, more trivial, extensions of the basic model are developed in an online appendix.[43]

---

[41]In general, the distribution of y may also depend on the actions of the short-run players.

[42]In the narrative of this paragraph, the reader should visualize the situation as follows. The long-run player of our model is the analyst. The commitment type is $\tau = G$, and the commitment action A is $\underline{q}_t$. The strategic type is $\tau = B$, and his deviating action D is $\underline{\sigma}_t \neq \underline{q}_t$. The short-run players of our model are the speculators. Their best response α to strategy A when the reputation of the analyst is good enough, is to blindly follow the recommendations of the analyst. So when reputation is high, (A, α) is a Nash equilibrium of the stage-game (this is Lemma 1): given that the bad analyst 'mimics' the good one then speculators want to follow blindly the recommendations; and given that screening breaks down the bad analyst then wants to 'imitate' the good one (this is Observation 1). As prices converge, α stops being a best response to A in our model and so, a fortiori, (A, α) does not remain a Nash equilibrium of the stage-game.

[43]This includes: adding public information about the asset's true value (e.g. dividend payments), allowing more general structures for the private information of investors, and adding other assets and/or analysts to the model.

5.1 Fundamental volatility and learning analyst ability

We assumed fixed fundamental values in the basic model. We now show that relaxing this assumption worsens considerably the market's ability to learn the true type of the analyst in the long run, expanding thereby the scope of our paper's main insights.[44]

The intuition is simple. With fixed fundamentals, whenever the market is trapped historical information embedded in prices eventually takes precedence over the current recommendation of the analyst. When this occurs, investors start turning to their private information with a view to evaluate the credence of this historical information. At that point, the market stops being trapped (c.f. Section 4). With volatile fundamentals by contrast, current recommendations may continue to be highly informative relative to historical information. This in turn, allows the analyst to sustain his grip on the market.

The next proposition considers the limit case, in which $\{\theta_{\iota(t)}\}_t$ are i.i.d. asset values with mean zero, and realizations in $\{-1, 1\}$. While evidently unrealistic, the model allows us to convey an essential insight in a simple way.[45] Let in what follows $\hat{\lambda} = \max\{\lambda_0, 1 - 2\frac{1-\pi}{\pi}\frac{1-\gamma}{\gamma}\}$.

Proposition 3. *With i.i.d. fundamental values, the market's ability to learn the true type of the analyst collapses, and reputation is (uniformly) bounded above. There exists $\varepsilon > 0$ such that, for all $\mathcal{H}_t$ and in any equilibrium, $\lambda_t < \hat{\lambda} + \varepsilon < 1$, and*

$$\mathbb{P}(\lim \lambda_t = 1 | \tau = G) = 0$$
$$\mathbb{P}(\lim \lambda_t = 0 | \tau = B) < 1.$$

Proposition 3 underscores the importance of the remarks made on the role played by prices in bringing about learning of the analyst's true type. When asset fundamentals are very volatile, prices have no time to accumulate historic information, and the 'value' of speculators' information about analyst ability remains constant. In that case, if and when the analyst gains sufficient reputation, speculators will follow his recommendations always, resulting in an indefinite trap.

---

[44]Introducing fundamental volatility effectively *prolongs* the influence of the short-run effects uncovered in Section 3.

[45]This simple model is studied, e.g., in Benabou and Laroque (1992).

5.2 New recommendations and prices' behavior

We endogenize in this subsection the decision of the analyst to make a recommendation or not, and explore the implications of our model regarding trading activity and prices' behavior over time. We show that trading activity peaks when the market is trapped, and that the occurrence of reputational traps can explain sudden and large price movements. Finally, we look at how our model is consistent with the empirical fact that previous recommendations may exert a contemporaneous effect on prices.

We have so far assumed that the good analyst receives each period a signal correlated with the asset's true value. This has forced the bad analyst to make recommendations each period in order to avoid losing his reputation. In practice, naturally, even a good analyst may in some periods fail to possess new information concerning the asset. Enhancing the model in this way opens up important new strategic considerations for the bad analyst, who may now decide in any given period whether or not to make a recommendation at all. Avoiding recommendations is attractive insofar as it minimizes information released about one's true type, but may in equilibrium arise suspicion regarding the ability of the analyst. We next explore the implications of this trade-off.

Formally, the basic model is modified as follows. We focus on a single asset, to save on notation. Each period the analyst now makes a choice $r_t \in \{-1, 0, 1\}$, where $r_t = 0$ is used to indicate the absence of a recommendation in period t. The good analyst is informed (i.e. makes a recommendation) on average half of the time, so that

$$\mathbb{P}(r_t = 0 | \tau = G) = 1/2$$
$$\mathbb{P}(r_t = \theta | \tau = G) = \phi/2,$$

where $\phi \in (1/2, 1)$. The rest of the model is as described in Section 2, except that now if no recommendation is made in period t, then trader t remains uninformed about the true type of the analyst.[46] It is precisely this feature which makes it attractive for the bad analyst to avoid making a recommendation.

We begin by showing that the rate at which the analyst makes recommendations varies

---

[46]To be sure, the absence of a recommendation may in itself be informative in equilibrium about the true type of the analyst. But s_t is unobserved unless $r_t \neq 0$. This assumption is justified by the fact that the quality of analyst reports is informative about analyst ability, and an analyst who does not make a recommendation needs not write a report either.

with time.[47] The intuition is as follows. So long as speculators screen the recommendations, the uninformed analyst will prefer to avoid making recommendations, so as to minimize the information released concerning his true ability. In a reputational trap however, the analyst knows that the market will follow his recommendations with probability one. He then needs not fear making new recommendations. The rate of recommendations thus rises when the market is trapped.

Observation 2. *The rate of recommendations varies with time, reaching a maximum when the market is trapped.*

We next record an important implication of the former observation. The arrival of new recommendations evidently provides speculators with incentives to trade based on the information provided by the analyst. Thus more recommendations typically spur more trade. The probability that a recommendation induces trade however, is normally less than one, unless namely the market finds itself in a reputational trap. We hence obtain:

Observation 3. *Trading activity varies with time, and peaks when the market is trapped.*

We turn next to the implications of our model regarding prices' behavior. Let $d_t := \left| \sum_0^{t-1} r_{t'} \right|$ record the 'net' amount of information concerning the asset accumulated up to period t.

Observation 4. *The greater the information accumulated about the asset, the more sensitive the price to the reputation of the analyst. Formally:*

$$\frac{\partial}{\partial d_t} \left(\left| \frac{\partial p_t^0}{\partial \lambda_t} \right| \right) > 0.$$

Observation 4 allows us to see why the occurrence of reputational traps can explain sudden and large price movements. As the market gets trapped it stops aggregating information about the analyst, but goes on accumulating recommendations stemming from the analyst (at an increased rate moreover, given Observation 2). All information accumulated *during* the reputational trap is then re-evaluated at once the instant the market exits the trap, and learning about the analyst's true type starts operating again. Interestingly, rather than being caused by the release of new information (as in, e.g., Lee (1998)), the large price 'swings' depicted here results from the *simultaneous depreciation of all past accumulated information.*

---

[47]Strictly speaking, only the rate at which the *bad* analyst makes new recommendations varies with time.

We complete this section with an important remark. Prices take time in the model we explore to fully incorporate the recommendations of the analyst. Here, a recommendation is first incorporated 'at a discount', and later adapted according to the evolution of the analyst's reputation.[48] Past recommendations thus affect prices contemporaneously, so long as the market goes on learning the true ability of the analyst.

Observation 5. *All past recommendations have a contemporaneous effect on the price, unless the market is trapped, in which case only the latest recommendation affects the price.*

5.3 Farsighted analyst

We modify here the basic model by relaxing the assumption that the analyst maximizes his expected reputation one period ahead, and assume instead that the uninformed analyst aims to maximize the expected discounted flow of his future income, where income in a given period is an increasing function of the analyst's reputation in that period. Let δ denote the discount factor and $f(\cdot)$ an increasing function relating income to reputation. Each period t, the analyst of type $\tau = B$ maximizes $\mathbb{E}_t[\sum_{s=t+1}^{\infty} \delta^{s-(t+1)} f(\lambda_s)|B]$, where $f'(\cdot) > 0$. The basic model is retrieved by setting $\delta = 0$.

The basic forces inducing screening to break down in the framework we explore are largely unaffected by the time horizon of the strategic analyst. As a sketch argument, suppose the uninformed analyst plays $\underline{\sigma}_t$ in period t, and let $m_t(\sigma_t^+) = \lambda_t q_t^+ + (1 - \lambda_t)\sigma_t^+$. Assuming as usual a bullish history, screening will break down on the positive side if

$$\frac{\lambda_t(1-\pi)q_t^+}{\lambda_t(1-\pi)q_t^+ + (1-\lambda_t)\sigma_t^+} \cdot \theta_t^G(1) \geq \frac{\gamma\lambda_t q_t^+}{\gamma m_t(\sigma_t^+) + 1 - \gamma} \cdot \theta_t^G(1) + \frac{1-\gamma}{\gamma m_t(\sigma_t^+) + 1 - \gamma} \cdot p_t^0.$$

Recall $q_t^+ \geq 1 - \phi > 0$. So if $\pi < 1$ and $\mu < 1$ (implying $\gamma < 1$) then as λ_t goes to one, the left-hand side goes to $\theta_t^G(1)$, and the right-hand side goes to some strictly convex combination of $\theta_t^G(1)$ and θ_t^G, which is obviously less than $\theta_t^G(1)$. Thus, for any σ_t^+ there exists a threshold for λ_t, above which the highlighted inequality will hold. A uniform threshold is then retrieved by letting σ_t^+ span $[0, 1]$.

Consider next the effect of the time horizon on the behavior of the strategic analyst. If screening were constant, i.e. if speculators' strategies were time-invariant, then an increase in today's reputation would induce a first-order stochastic dominance shift of reputation in

---

[48]See, e.g., Michaely and Womack (2005) for empirical evidence of such behavior.

any future period. The objective of the analyst in that case would reduce to maximizing reputation one period ahead. In the framework we explore, however, the degree of screening carried out by speculators varies with time. The time horizon will therefore in principle affect the behavior of the strategic analyst, who may be willing to sacrifice reputation today in order to reap the benefits from less screening some time in the future.

We now illustrate (informally first, then with a proposition) why these effects are unimportant for the qualitative results of our paper. Suppose that the analyst looks forward n periods, i.e. the analyst cares about his stream of income (and hence, his reputation) starting from period $t+1$ all the way to period $t+n$. Define $\hat{\lambda}_k(\mathcal{R}_t)$, $k \in \{1, ..., n\}$, as follows. First, set $\hat{\lambda}_1(\mathcal{R}_t) = \hat{\lambda}(\mathcal{R}_t)$, the threshold elicited in Lemma 1 of Section 3.[49] Set then, recursively, $\hat{\lambda}_k(\mathcal{R}_t) = \max\{\hat{\lambda}_{k-1}(\mathcal{R}_t, +1), \hat{\lambda}_{k-1}(\mathcal{R}_t, -1)\}$. The idea is the following. If $\lambda_t \geq \hat{\lambda}_n(\mathcal{R}_t)$ then screening will break down in at least the next n periods. But the analyst cares about n periods only. So if today his reputation is above $\hat{\lambda}_n(\mathcal{R}_t)$ then effectively, from his perspective, the strategy of speculators *is* time-invariant. This simple example thus suggests how, by raising appropriately the threshold reputation of the analyst, results and insights obtained in the basic model do carry over as we extend the time horizon of the strategic analyst.

Finally, we summarize the previous remarks in a formal proposition, and show in addition that the long-run behavior of the basic model is unaffected by the time horizon of the analyst.

Proposition 4. *Let $\delta \in [0, 1)$, $\mu < 1$, and $\pi < 1$. Fix $\epsilon > 0$. There exists a threshold reputation $\hat{\lambda}(\mathcal{H}_t, \delta)$ such that, in any equilibrium with $\lambda_t \geq \hat{\lambda}(\mathcal{H}_t, \delta)$:*

$$\mathbb{P}_t\big(|\lambda_{t+1} - \lambda_t| < \epsilon\big) = 1,$$

given that history. As δ tends to zero, $\hat{\lambda}(\mathcal{H}_t, \delta) \to \hat{\lambda}(\mathcal{H}_t)$, the threshold reputation elicited in Lemma 1 of Section 3. Furthermore, in the long run, the market learns τ almost surely in any equilibrium. In particular:

1. *Conditional on type $\tau = G$, prices converge almost surely to the true asset values.*

2. *Conditional on type $\tau = B$, prices converge almost surely to zero.*

Results obtained in the basic model thus approximate well those of the more general framework considered here.

---

[49]Notice that the thresholds from Section 3 are functions of $\mathcal{H}_t$ through $\mathcal{R}_{t-1}$ and $\mathcal{I}_t$ only. For convenience, we here suppose that there is only one asset and define $\hat{\lambda}(\mathcal{R}_t) := \hat{\lambda}(\mathcal{H}_{t+1})$.

6 Conclusion

Do asset prices aggregate investors' private information concerning financial analysts? We have uncovered in this paper a market failure previously ignored, by showing that investors' private information concerning financial analysts is typically not aggregated into prices in the medium run. This phenomenon – which we coin *reputational trap* – could allow incompetent analysts to sustain a good reputation longer than they should, and implies that bad analysts will tend to cause substantial short-term price fluctuations: in a bullish market, a bad analyst drives prices up when the market is trapped, inducing sudden price reversals as (and if) the market eventually learns the true type of the analyst.

We also present several results with regards to the long-run convergence of prices. The difficulty lies in the manner in which the market provides feedback about analyst ability. Effectively, due to the existence of reputational traps, the market's feedback about analysts is not 'bounded away from zero'. It turns out that the driver of convergence in our model is the price mechanism itself, which ensures that the more information becomes released about an asset, the more profitable for speculators to screen analysts. However, when assets' fundamentals are very volatile then historic information rapidly depreciates. The 'value' of speculators' information about analysts remains roughly constant, in that case. Such conditions may induce perpetual, or very long, reputational traps.

This paper offers several avenues for future research, two of which we now discuss. First, we have assumed that analysts make recommendations on one asset per period, and have restricted speculators to trading precisely this asset. If we relaxed the choice of asset traded in period t, interesting effects would ensue. Consider for instance a speculator with $s_t = 1$ receiving a recommendation regarding asset $\iota(t)$. He would then revalue not just asset $\iota(t)$, but all assets at once, since he now realizes that all previous recommendations were indeed 'true'. Thus, the best trade he can make may not be in asset $\iota(t)$, but in some other under/overvalued asset. Second, we have assumed that the analyst is motivated by career concerns. It would be interesting to know whether and how the conclusions of our paper would change if the incentives of the analyst were commissions-based instead, i.e. if the objective of the analyst were to maximize the trading activity.

7 Appendix

In order to avoid repeating very similar proofs twice, we will work in this appendix with the slightly more general model of Section 5.2, i.e. we allow the analyst not to make a recommendation in a given period.

To shorten notation, we will make use throughout the appendix of

$$\gamma := \frac{\mu}{\mu + (1-\mu)/3},$$

$$p_t^0 = \mathbb{E}_t[\theta],$$

$$\lambda_{t+1}^e(r, \underline{\sigma}_t, \underline{\xi}_t) := \mathbb{E}_t[\lambda_{t+1}(r_t = r, y_t)|\tau = B; \underline{\sigma}_t, \underline{\xi}_t].$$

The parameter γ records the weight of speculators relative to liquidity traders taking any given action (buy/sell/abstain). The variable p_t^0 is the public valuation at the beginning of period t. The variable $\lambda_{t+1}^e(r, \underline{\sigma}_t, \underline{\xi}_t)$ finally, denotes the uninformed analyst's expected reputation next period from recommending r today, when his strategy is $\underline{\sigma}_t$ and speculators' strategy is $\underline{\xi}_t$.

We will, wherever possible, drop time subscripts in order to unclutter notation. The following equilibrium conditions, stating that the uninformed analyst is indifferent between all recommendations, are easily established using standard arguments:

$$\lambda^e(-1, \underline{\sigma}, \underline{\xi}) = \lambda^e(0, \underline{\sigma}, \underline{\xi}) = \lambda^e(1, \underline{\sigma}, \underline{\xi}). \tag{8}$$

7.1 Proofs of Lemma 2 and Theorem 1

We begin with a series of preliminary results. Claims 1-3 establish the intuitive results that more screening from speculators and more aggressive behavior from the uninformed analyst induce worse expected reputation for the analyst. Claim 4 orders speculators' valuations, in equilibrium. Claims 5 and 6 are instrumental in proving equilibrium uniqueness.

Claim 1. *Consider $\underline{\sigma}$ such that $\sigma^r > 0$, $\forall r$.*

1. *If $\underline{\xi}_1$ entails a break-down of screening on the positive side then $\lambda^e(1, \underline{\sigma}, \underline{\xi}_1) \geq \lambda^e(1, \underline{\sigma}, \underline{\xi}_2)$, with strict inequality unless $\underline{\xi}_2$ entails a break-down of screening on the positive side too.*

2. *If $\underline{\xi}_1$ entails a break-down of screening on the negative side then $\lambda^e(-1, \underline{\sigma}, \underline{\xi}_1) \geq \lambda^e(-1, \underline{\sigma}, \underline{\xi}_2)$, with strict inequality unless $\underline{\xi}_2$ entails a break-down of screening on the negative side too.*

Proof of Claim 1: Let, for $\mathbb{P}(y|r,\underline{\xi},B) > 0$:

$$L(y|r,\underline{\xi}) = \frac{\mathbb{P}(y|r,\underline{\xi},G)}{\mathbb{P}(y|r,\underline{\xi},B)}.$$

Note that

$$\mathbb{E}[L(y|r,\underline{\xi})|r,\underline{\xi},B] = \sum_{y:\mathbb{P}(y|r,\underline{\xi},B)>0} \mathbb{P}(y|r,\underline{\xi},B)\frac{\mathbb{P}(y|r,\underline{\xi},G)}{\mathbb{P}(y|r,\underline{\xi},B)} = 1.$$

Using Bayes' rule:

$$\lambda^e(r,\underline{\sigma},\underline{\xi}) = \mathbb{E}[M_{\sigma^r}(L(y|r,\underline{\xi}))|r,\underline{\xi},B], \tag{9}$$

where $M_{\sigma^r}(x) = \frac{\lambda q^r x}{\lambda q^r x + (1-\lambda)\sigma^r}$ is concave.

Observe next that if $\underline{\xi}_1$ entails a break-down of screening on the positive side then $L(y|r = 1,\underline{\xi}_1) = 1$ for all y. So either $\underline{\xi}_2$ entails a break-down of screening on the positive side too or, conditional on the expert being bad, the distribution of $L(y|r = 1,\underline{\xi}_2)$ is a mean-preserving spread of the distribution of $L(y|r = 1,\underline{\xi}_1)$. In the latter case we obtain, by (9) and concavity of M_{σ^+}, $\lambda^e(1,\underline{\sigma},\underline{\xi}_2) < \lambda^e(1,\underline{\sigma},\underline{\xi}_1)$.

The proof of Part 2 of the Claim is similar, and omitted.

∎

Claim 2. *If $\underline{\xi}_1$ entails a break-down of screening then for any $\underline{\xi}_2$:*

1. $\sigma_2^+ > \sigma_1^+ \Rightarrow \lambda^e(1,\underline{\sigma}_1,\underline{\xi}_1) > \lambda^e(1,\underline{\sigma}_2,\underline{\xi}_2)$.

2. $\sigma_2^- > \sigma_1^- \Rightarrow \lambda^e(-1,\underline{\sigma}_1,\underline{\xi}_1) > \lambda^e(-1,\underline{\sigma}_2,\underline{\xi}_2)$.

Proof of Claim 2: Using the same notation as in the proof of Claim 1:

$$\begin{aligned}
\lambda^e(1,\underline{\sigma}_2,\underline{\xi}_2) &= \mathbb{E}[M_{\sigma_2^+}\big(L(y|r = 1,\underline{\xi}_2)\big)|r = 1,\underline{\xi}_2,B] \\
&< \mathbb{E}[M_{\sigma_1^+}\big(L(y|r = 1,\underline{\xi}_2)\big)|r = 1,\underline{\xi}_2,B] \\
&= \lambda^e(1,\underline{\sigma}_1,\underline{\xi}_2) \\
&\leq \lambda^e(1,\underline{\sigma}_1,\underline{\xi}_1).
\end{aligned}$$

The first inequality follows from the fact that $M_{\sigma_2^+}(x) < M_{\sigma_1^+}(x)$ for all x. The last inequality is an application of Claim 1.

28

The proof of Part 2 of the Claim is similar, and omitted.

∎

Claim 3. *Let $r, r' \in \{-1, 1\}$. If for $y \neq \hat{y}$ we have $\xi^y(r, 0) = \xi^{\hat{y}}(r, 1) = 1$ and moreover $\frac{\sigma^r}{q^r} > \frac{\sigma^{r'}}{q^{r'}}$, then $\lambda^e(r', \underline{\sigma}, \underline{\xi}) > \lambda^e(r, \underline{\sigma}, \underline{\xi})$.*

Proof of Claim 3: The proof is similar to that of Claim 2 and is therefore omitted.

∎

We next prove equations (3) and (4) of Section 3. Recall that $v(r, s)$ denotes the valuation of an investor with recommendation r and ability signal s. By extension, $v(0)$ is used to denote the valuation of an investor without a recommendation.[50]

---

[50]Recall in particular that under the assumptions of Section 5.2, s_t is unobserved unless $r_t \neq 0$.

Claim 4. *Let $\theta^G(-1) > 0$. Then, in any equilibrium:*

1. *$v(1,1)$ is (strictly) the highest of all valuations.*

2. *$v(-1,0)$ is (strictly) the lowest of all valuations.*

3. *$v(-1,0) < p^0 < v_t(1,1)$.*

In particular, in any equilibrium: $\xi^a(1,1) = 1$ and $\xi^b(-1,0) = 1$.

Proof of Claim 4: Part 1 is immediate. We prove the second part. Let $(\underline{\sigma}, \underline{\xi})$ denote an arbitrary equilibrium strategy-pair.

Step 1: $v(-1,0) < v(0)$. Let $\beta(r,s)$ denote an investor's updated belief of the analyst's type after observing (r,s). We have:

$$\lambda^e(0, \underline{\sigma}, \underline{\xi}) = \lambda^e(-1, \underline{\sigma}, \underline{\xi}) > \beta(-1,0).$$

Step 1 now follows, since:

$$v(-1,0) = \beta(-1,0)\theta^G(-1) < \lambda^e(0,\underline{\sigma},\underline{\xi})\theta^G(-1) = \beta(0)\theta^G(-1) < \beta(0)\theta^G(0) = v(0).$$

Step 2: $v(-1,0) < v(1,0)$. Suppose, for the sake of contradiction, that $v(-1,0) \geq v(1,0)$. Then $\beta(-1,0) > \beta(1,0)$, and so $\sigma^-/q^- < \sigma^+/q^+$. But $v(-1,0) \geq v(1,0)$ also implies (using Step 1) that $v(1,0)$ is the lowest valuation, in which case $\xi^b(1,0) = 1$. Since $\xi^a(1,1) = 1$ (by Part 1 of the Claim), then Claim 3 yields $\lambda^e(1, \underline{\sigma}, \underline{\xi}) < \lambda^e(-1, \underline{\sigma}, \underline{\xi})$. But this is impossible, in equilibrium.

To prove part 3 of the claim, note that $v(-1,0) \geq p^0$ implies $\beta(-1,0) > \lambda$. But then $\lambda^e(-1, \underline{\sigma}, \underline{\xi}) > \lambda$, which is impossible by (8). That $p^0 < v(1,1)$ is immediate.

∎

Our next result establishes a key step in the proof of Lemma 2.

Claim 5. *Let $\theta^G(-1) > 0$, and E_i, $i = 1, 2$, two equilibria. Let $\Delta\sigma^r := \sigma_2^r - \sigma_1^r$. If $\Delta\sigma^- < 0$, and $\Delta\sigma^0, \Delta\sigma^+ \geq 0$, then $p_1^b \leq p_2^b$.*

Proof of Claim 5: We give here a very general proof of the result. A less abstract (but longer) proof can be found in Rudiger and Vigier (2014). The proof we give requires some notation and terminology, which we now define. Let $I := \{0, ..., n\}$. Consider a discrete random variable θ, with realizations $\{\theta_i\}_{i \in I}$, $\theta_i < \theta_j$ for $i < j$, and probability distribution $G_i := \mathbb{P}(\theta = \theta_i)$, $i \in I$. Define a signal s by the non-negative weights $\{W_i(s)\}_{i \in I}$, where $W_i(s) \leq G_i$, $\forall i$. Let moreover $W(s) := \sum_I W_i(s)$, and $E(s) := \sum_I \frac{W_i(s)}{W(s)} \theta_i$. Given two signals s and s' with $W_i(s) + W_i(s') \leq G_i$, $i \in I$, define a new signal $s + s' := \{W_i(s) + W_i(s')\}_{i \in I}$. Similarly, given two signals s and s' with $W_i(s') \leq W_i(s)$, $i \in I$, define a new signal $s - s' := \{W_i(s) - W_i(s')\}_{i \in I}$. The following property follows directly: $E(s + s') = \frac{W(s)}{W(s) + W(s')} E(s) + \frac{W(s')}{W(s) + W(s')} E(s')$.

The interpretation should be clear. A signal s is obtained by taking probability mass from the different realizations of the random variable θ; $W(s)$ is the total probability that s obtains given the prior on θ, and $E(s)$ is the posterior mean of θ when s is observed. We now make the link with the model of this paper. At the beginning of an arbitrary period where $\theta^G(-1) > 0$, the asset's true value θ can be viewed as a discrete random variable with (ordered) realizations $\{0, \theta^G(-1), \theta^G(0), \theta^G(-1)\}$, and probability distribution $G_0 = 1 - \lambda$, $G_1 = \lambda q^-$, $G_2 = \lambda q^0$, and $G_3 = \lambda q^+$. The probability mass is then split into 6 different signals: s_a, corresponding to the information of speculators with $(r, s) = (-1, 0)$, s_b corresponding to the information of speculators with $(r, s) = (-1, 1)$, s_c corresponding to $r = 0$, s_d corresponding to $(r, s) = (1, 0)$, s_e corresponding to $(r, s) = (1, 1)$, and $\tilde{s}$ corresponding to the information of liquidity traders. The signals are easily computed as functions of $\underline{\sigma}$ (and of the primitives of the model). The signal s_a, e.g., is given by $W_0(s_a) = \mu \sigma^-$, $W_1(s_a) = \mu(1 - \pi)q^-$, and $W_2(s_a) = W_3(s_a) = 0$. The signal $\tilde{s}$ is obtained by retrieving probability mass uniformly from the prior distribution. Hence: $W_i(\tilde{s}) = (1 - \mu)G_i$, $\forall i$.

Now fix a candidate bid price $z \in [-1, 1]$. For any z, we can define a new signal $\hat{s}_z$ as

$$\hat{s}_z := \frac{1}{3}\tilde{s} + \sum_{E(s_k) < z} s_k.$$

Intuitively, the signal $\hat{s}_z$ corresponds to the information available given that a sell order was passed at bid price $p^b = z$. Note in particular that the unique equilibrium bid price – for $\underline{\sigma}$ fixed – is given by $\inf\{z : E(\hat{s}_z) \leq z\}$.

We can now conclude the proof of the claim. Suppose there exist two equilibria, E and E' say, where E entails $\underline{\sigma}$ and E' entails $\underline{\sigma}'$, and which satisfy the conditions laid out in the

statement of the claim, i.e. such that $\underline{\sigma}'$ may be obtained from $\underline{\sigma}$ by shifting weight away from $r = -1$ and onto other recommendations. Observe that, for any $z \in [-1, 1]$, we can decompose

$$\hat{s}'_z = (\hat{s}_z - \underline{s}) + \overline{s}_z,$$

where $W_0(\underline{s}) = W(\underline{s}) = \sigma^- - \sigma'^-$, and $E(\overline{s}_z) \geq z$. The intuition is the following. As we move from $\underline{\sigma}$ to $\underline{\sigma}'$ we shift 'bad' probability mass from the signal with the worse posterior, s_a, to signals with better posteriors, s_c and s_d.[51] Either $E(s'_c)$ and/or $E(s'_d)$ fall below z, or they don't. If they don't then $\hat{s}'_z$ is obtained from $\hat{s}_z$ by simply removing some probability mass originating from G_0. If they do then $\hat{s}'_z$ is obtained from $\hat{s}_z$ by adding s_c and s_d; but in that case, by construction: $E(s_c), E(s_d) \geq z$.

Observe next that

$$E(\hat{s}'_z) \geq E(\hat{s}_z + \overline{s}_z) \geq \frac{W(\hat{s}_z)}{W(\hat{s}_z) + W(\overline{s}_z)} E(\hat{s}_z) + \frac{W(\overline{s}_z)}{W(\hat{s}_z) + W(\overline{s}_z)} z,$$

and hence:

$$E(\hat{s}_z) \geq z \Rightarrow E(\hat{s}'_z) \geq z.$$

We obtain finally, using an earlier remark, that $p^{b'} \geq p^b$, and the claim is established.

∎

The next claim was proven within the proof of Claim 5. We state it here for the record. In what follows, for a given equilibrium, let $\underline{p}$ denote the vector of bid and ask prices given by (2).

Claim 6. *Prices are uniquely determined by the strategy of the uninformed analyst. Let E_i, $i = 1, 2$ denote two equilibria. Then:*

$$\underline{\sigma}_1 = \underline{\sigma}_2 \Rightarrow \underline{p}_1 = \underline{p}_2.$$

Proof of Lemma 2: As usual we work the proof for the case where $\theta_t^G(-1) > 0$. Other cases can be treated similarly. Let E_i, $i = 1, 2$ denote two equilibria, and assume that in E_1 screening breaks down. The proof has four steps.

---

[51]C.f. Claim 4.

Step 1: $\sigma_2^+ \geq \sigma_1^+$. Step 1 is proved by contradiction. Suppose $\sigma_2^+ < \sigma_1^+$. Then $\xi_2^a(1,0) = \xi_1^a(1,0) = 1$. Hence, from Claim 2: $\lambda^e(1, \underline{\sigma}_2, \underline{\xi}_2) > \lambda^e(1, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$. Now either $\sigma_2^0 > \sigma_1^0$ (call this Case 1), or $\sigma_2^- > \sigma_1^-$ (call this Case 2). Case 1 immediately gives $\lambda^e(0, \underline{\sigma}_2, \underline{\xi}_2) < \lambda^e(0, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$. Case 2 gives $\lambda^e(-1, \underline{\sigma}_2, \underline{\xi}_2) < \lambda^e(-1, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$, where we have made use again of Claim 2. Either way, we obtain a contradiction with (8) in the second equilibrium.

Step 2: $\sigma_2^- \geq \sigma_1^-$. The proof of Step 2 is again by contradiction. Suppose $\sigma_2^- < \sigma_1^-$. Three cases must be considered: $\Delta\sigma^0 \geq 0$ and $\Delta\sigma^+ \geq 0$ (Case 1), $\Delta\sigma^0 > 0$ and $\Delta\sigma^+ < 0$ (Case 2), $\Delta\sigma^0 < 0$ and $\Delta\sigma^+ > 0$ (Case 3). In Case 1, Claim 5 yields $p_2^b \geq p_1^b$. Hence (generically) $\xi_2^b(-1,1) = \xi_1^b(-1,1) = 1$, from which we obtain using Claim 2 that $\lambda^e(-1, \underline{\sigma}_2, \underline{\xi}_2) > \lambda^e(-1, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$. But then we have a contradiction with (8) in the second equilibrium since $\lambda^e(0, \underline{\sigma}_2, \underline{\xi}_2) \leq \lambda^e(0, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$. Consider Case 2 next. Reproducing the arguments of Step 1 yields $\lambda^e(0, \underline{\sigma}_2, \underline{\xi}_2) < \lambda^e(0, \underline{\sigma}_1, \underline{\xi}_1) = \lambda = \lambda^e(1, \underline{\sigma}_1, \underline{\xi}_1) < \lambda^e(1, \underline{\sigma}_2, \underline{\xi}_2)$. Again, we obtain a contradiction with (8) in the second equilibrium. Consider finally Case 3. Using Claim 2 yields $\lambda^e(0, \underline{\sigma}_2, \underline{\xi}_2) > \lambda^e(0, \underline{\sigma}_1, \underline{\xi}_1) = \lambda = \lambda^e(1, \underline{\sigma}_1, \underline{\xi}_1) > \lambda^e(1, \underline{\sigma}_2, \underline{\xi}_2)$. Yet again, we obtain a contradiction with (8) in the second equilibrium.

Step 3: $\underline{\sigma}_2 = \underline{\sigma}_1$. If not, then combining Steps 1 and 2 yields $\sigma_2^0 < \sigma_1^0$ and hence $\lambda^e(0, \underline{\sigma}_2, \underline{\xi}_2) > \lambda^e(0, \underline{\sigma}_1, \underline{\xi}_1) = \lambda$. This necessarily contradicts (8) in the second equilibrium, since the bad analyst is unable to improve his reputation on average.

Step 4: $E_2 = E_1$. Combining Step 3 and Claim 6 yields $\underline{p}_1 = \underline{p}_2$. So $\underline{\xi}_1 = \underline{\xi}_2$ (generically). ∎

Proof of Theorem 1: Combining Lemmas 1 and 2 gives for each history $\mathcal{H}_t$ a minimum threshold reputation such that if at time t, given that history, λ_t is above the threshold, then the market must be trapped (in any equilibrium). Note moreover that the threshold $\hat{\lambda}(\mathcal{H}_t)$ depends on history $\mathcal{H}_t$ only through $\mathcal{I}_t$ and the recommendations in $\mathcal{R}_{t-1}$. For fixed T therefore, the set of possible contemporaneous thresholds at time T is finite. Inspection of the expressions given in the text for $\hat{\lambda}^-(\mathcal{H}_t)$ and $\hat{\lambda}^+(\mathcal{H}_t)$ shows that the maximum threshold is attained at histories such that the same recommendation was made on the same asset T times. Setting $\hat{\lambda}_0^T$ equal to the maximum threshold at T completes the proof of the first part

of the theorem. The second part of the theorem then follows immediately from taking the relevant derivatives.

∎

7.2   Proofs of Proposition 2 and Theorem 2

Before proving Proposition 2 and Theorem 2, we establish a series of claims. Claim 7 establishes a lower bound on the rate at which the uninformed analyst issues recommendations each period. Claim 8 lower-bounds the probability that a significant jump in reputation occurs, whenever screening is efficient. Claim 9 shows that unless reputation converges to zero or one then screening will eventually become efficient. Claim 10 concludes that reputation converges to either zero or one, almost surely. Finally, Claim 11 shows that reputation converges to the 'correct' value. With these results in hand, we then prove the theorem.

Claim 7. *There exists $\delta > 0$ such that in any equilibrium, for any history $\mathcal{H}_t$ and recommendation r: $\sigma_t^r > \delta$.*

Proof of Claim 7: Notice that due to liquidity traders we can find $\ell > 0$ such that for any r_t, y_t and $\underline{\xi}_t$ then (using notation from Claim 1):

$$\ell^{-1} < L_t(y_t | r_t, \underline{\xi}_t) < \ell.$$

Note also that for any history $\mathcal{H}_t$ and any recommendation r then $q_t^r \in [\frac{1-\phi}{2}, \frac{1}{2}]$. Thus for any $\mathcal{H}_t$, r_t, y_t and $\underline{\xi}_t$:

$$\frac{\mathbb{P}_t(y_t, r_t | \underline{\sigma}_t, \underline{\xi}_t, G)}{\mathbb{P}_t(y_t, r_t | \underline{\sigma}_t, \underline{\xi}_t, B)} = \frac{q_t^{r_t} \mathbb{P}_t(y_t | r_t, \underline{\xi}_t, G)}{\sigma_t^{r_t} \mathbb{P}_t(y_t | r_t, \underline{\xi}_t, B)} > \frac{(1-\phi)\ell^{-1}}{2\sigma_t^{r_t}}.$$

Hence:

$$\sigma_t^{r_t} < \frac{1-\phi}{2}\ell^{-1} \Rightarrow \frac{\mathbb{P}_t(y_t, r_t | \underline{\sigma}_t, \underline{\xi}_t, G)}{\mathbb{P}_t(y_t, r_t | \underline{\sigma}_t, \underline{\xi}_t, B)} > 1.$$

The claim now follows by setting $\delta = \frac{1-\phi}{2}\ell^{-1}$, since if the bad analyst were ever to issue recommendation r with less than probability δ, then he could, by recommending r, improve his reputation with certainty. But this is impossible in equilibrium.

∎

34

Claim 8. *For all $\tilde{\lambda} \in (0,1)$, there exists $\varepsilon_{\tilde{\lambda}} > 0$ and $\delta_{\tilde{\lambda}} > 0$ with the following property: if $|\lambda_t - \tilde{\lambda}| < \delta_{\tilde{\lambda}}$ then, in any equilibrium where screening is efficient during period t:*

$$\mathbb{P}_t(|\lambda_{t+1} - \lambda_t| \geq \varepsilon_{\tilde{\lambda}}) \geq b,$$

where $b > 0$ is independent of $\tilde{\lambda}$.

Proof of Claim 8: Let $\tilde{\lambda} \in (0,1)$, $\lambda_t = \tilde{\lambda}$, and consider an equilibrium where screening is efficient. We will show that we can find $\varepsilon_{\tilde{\lambda}}$, and b independent of $\tilde{\lambda}$ such that the highlighted equation of the claim holds.

One recommendation $r^* \in \{-1, 1\}$ must have $\sigma_t^{r^*} \leq q_t^{r^*}$ (this follows from Observation 2). Let y^* denote the action of the speculator with $(r_t = r^*, s_t = 1)$. Since screening is efficient: $\xi_t^{y^*}(r^*, 0) = 0$. Using (1) yields

$$
\begin{aligned}
\lambda_{t+1}(r_t = r^*, y_t = y^*) &= \frac{\lambda_t q_t^{r^*}(\gamma\pi + (1-\gamma))}{\lambda_t q_t^{r^*}(\gamma\pi + (1-\gamma)) + (1-\lambda_t)\sigma_t^{r^*}(1-\gamma)} \\
&\geq \frac{\lambda_t(\gamma\pi + (1-\gamma))}{\lambda_t(\gamma\pi + (1-\gamma)) + (1-\lambda_t)(1-\gamma)} \\
&> \lambda_t.
\end{aligned}
$$

Let $\varepsilon_{\tilde{\lambda}}$ be defined by the difference between the final two terms of the right-hand side:

$$\varepsilon_{\tilde{\lambda}} = \frac{\tilde{\lambda}(\gamma\pi + (1-\gamma))}{\tilde{\lambda}(\gamma\pi + (1-\gamma)) + (1-\tilde{\lambda})(1-\gamma)} - \tilde{\lambda}.$$

Then $\varepsilon_{\tilde{\lambda}}$ and $b = \frac{1-\mu}{3} \cdot \min\{\delta, \frac{1-\phi}{2}\}$ together satisfy the conditions we were looking for, with δ defined by Claim 7.

The claim follows by continuity of $\varepsilon_{\tilde{\lambda}}$ (defined above) as a function of $\tilde{\lambda}$. ∎

The sequence $\{\lambda_t\}$ is a bounded martingales with respect to the filtration $\{\mathcal{H}_t\}$. The Martingale Convergence Theorem thus applies. We let in what follows the random variable λ denote the (a.s.) limit of the sequence $\{\lambda_t\}$.

Claim 9. *Let W denote the event $\lambda \notin \{0,1\}$. Then, conditional on W, screening eventually becomes efficient, a.s..*

Proof of Claim 9: Focus here on a single asset, in order to save on notation (the proof extends in the obvious way to any *finite* number of assets). The process $\{\theta_t^G(0)\}$ is a bounded martingale with respect to the filtration $\{\mathcal{H}_t\}$, under $\tau = G$. The Martingale Convergence Theorem thus applies. Moreover, clearly, $\lim_{t\to\infty} \theta_t^G(0) \in \{-1,1\}$. This implies, in turn, that either $\lambda = 0$, or $\lim_{t\to\infty} \theta_t^G(0)$ exists and takes values in $\{-1,1\}$. The random variable $\lim_{t\to\infty} \theta_t^G(0)$ is thus well-defined under W. Denote $\theta^G := \lim_{t\to\infty} \theta_t^G(0)$, where the random variable is now defined over the entire W.

We condition henceforth on W. Suppose to fix ideas that $\theta^G = 1$ (the case $\theta^G = -1$ is similar and omitted). Then, for any $\epsilon > 0$, we can find a time from which point onwards $|\theta_t^G(r) - 1| < \epsilon$, for all r. Choosing ϵ small enough yields from some time onwards $v_t(-1,1) = \theta_t^G(-1) > \lambda_t \theta_t^G = p_t^0$, since by assumption the limit $\lambda < 1$. Next, applying Claim 4 gives

$$v_t(-1,0) < p_t^b < p_t^0 < v_t(-1,1),$$

and shows that screening eventually becomes efficient (a.s.) on the negative side.

We proceed to show that screening eventually becomes efficient on the positive side too. Applying again Claim 4, observe that

$$p_t^a \geq \frac{1-\gamma}{(1-\gamma) + \gamma \lambda_t \pi} p_t^0 + \frac{\gamma \lambda_t \pi}{(1-\gamma) + \gamma \lambda_t \pi} \theta_t^G(1).$$

As $\lambda > 0$, from some time onwards p_t^a is thus bounded away from p_t^0. But then $v_t(1,0) \geq p_t^a$ must imply $\beta_t(1,0) > \lambda_t$ (recall that $\beta(r,s)$ denotes an investor's updated belief of the analyst's type after observing (r,s)), which in turn implies $\lambda^e(1, \underline{\sigma}_t, \underline{\xi}_t) > \lambda_t$, contradicting (8). Hence $v_t(1,0) < p_t^a$, and screening eventually becomes efficient on the positive side too.

$\blacksquare$

Claim 10. *Let $\lambda = \lim_{t\to\infty} \lambda_t$. Then $\lambda \in \{0,1\}$, a.s..*

Proof of Claim 10: Let $W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2)$ denote the event $|\lambda - \tilde{\lambda}| < \delta_{\tilde{\lambda}}/2$, where $\tilde{\lambda} \in (0,1)$ and $\delta_{\tilde{\lambda}}$ as defined in Claim 8. Clearly, we can choose $\delta_{\tilde{\lambda}}$ such that $W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2) \subset W$, where as before W denotes the event $\lambda \in (0,1)$. Using Claim 9, for any $\omega \in W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2)$ we can (almost surely) define a smallest time $T(\omega)$ such that for all $s \geq T(\omega)$: (a) screening is efficient and (b) $|\lambda_s - \lambda| < \delta_{\tilde{\lambda}}/2$. Let $V_k = \{\omega \in W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2) : T(\omega) = k\}$. Then $W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2) = \bigcup V_k$.

Furthermore, applying Claim 8:

$$\mathbb{P}\big(|\lambda_{s+1} - \lambda_s| \geq \varepsilon_{\tilde{\lambda}}|V_k\big) \geq b \ , \quad s \geq k.$$

But $\{\lambda_t\}$ converges a.s., and hence also in probability. Hence, $\mathbb{P}\big(V_k\big) = 0$, for all k, and ultimately $\mathbb{P}\big(W(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2)\big) = 0$.

Let $B(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2)$ denote the open ball with center $\tilde{\lambda}$ and radius $\delta_{\tilde{\lambda}}/2$. For any n, the interval $[1/n, 1 - 1/n]$ has an open cover consisting of open balls $B(\tilde{\lambda}, \delta_{\tilde{\lambda}}/2)$, $\tilde{\lambda} \in [1/n, 1 - 1/n]$. By compactness, we can extract a finite sub-cover $\Big(B(\tilde{\lambda}_s, \delta_{\tilde{\lambda}_s}/2)\Big)_{s=1}^{S}$. Then

$$\mathbb{P}(\lambda \in [1/n, 1 - 1/n]) \leq \mathbb{P}\Big(\lambda \in \bigcup B(\tilde{\lambda}_s, \delta_{\tilde{\lambda}_s}/2)\Big) \leq \sum \mathbb{P}\Big(\lambda \in B(\tilde{\lambda}_s, \delta_{\tilde{\lambda}_s}/2)\Big) = 0.$$

This being true for all n, we finally obtain $\mathbb{P}(\lambda \in \{0, 1\}) = 1$.

∎

Claim 11. *The market learns τ, almost surely. Let $\lambda = \lim_{t \to \infty} \lambda_t$: if $\tau = G$ then $\lambda = 1$ a.s.; if $\tau = B$ then $\lambda = 0$ a.s..*

Proof of Claim 11: Consider first $\tau = B$. In any equilibrium:

$$\begin{aligned}
\mathbb{E}_t\left[\frac{\lambda_{t+1}}{1 - \lambda_{t+1}}\Big|B\right] &= \sum_{r_{t+1}, y_{t+1}} \mathbb{P}_t(r_{t+1}, y_{t+1}|B) \cdot \frac{\lambda_t \mathbb{P}_t\big(r_{t+1}, y_{t+1}|G\big)}{(1 - \lambda_t)\mathbb{P}_t(r_{t+1}, y_{t+1}|B)} \\
&= \frac{\lambda_t}{1 - \lambda_t} \sum_{r_{t+1}, y_{t+1}} \mathbb{P}_t(r_{t+1}, y_{t+1}|G) \\
&= \frac{\lambda_t}{1 - \lambda_t}.
\end{aligned}$$

Hence $\mathbb{E}\left[\frac{\lambda_t}{1-\lambda_t}\Big|B\right] = \frac{\lambda_0}{1-\lambda_0}$, for all t. Fatou's Lemma then gives $\mathbb{E}\left[\frac{\lambda}{1-\lambda}\Big|B\right] \leq \frac{\lambda_0}{1-\lambda_0}$. Hence $\mathbb{P}(\lambda = 1|B) = 0$.

Similar derivations give $\mathbb{E}\left[\frac{1-\lambda_t}{\lambda_t}\Big|G\right] = \frac{1-\lambda_0}{\lambda_0}$, for all t, and so $\mathbb{E}\left[\frac{1-\lambda}{\lambda}\Big|G\right] \leq \frac{1-\lambda_0}{\lambda_0}$. Hence $\mathbb{P}(\lambda = 0|G) = 0$.

Claim 11 now follows immediately, by application of Claim 10.

∎

Proof of Proposition 2: Fix an arbitrary history $\mathcal{H}_t$ with $\iota(t) = i'$. If the analyst has reputation $\lambda_t \geq \hat{\lambda}(\mathcal{H}_t)$ we are done. If he has reputation $\lambda_t < \hat{\lambda}(\mathcal{H}_t)$, construct a new history $\mathcal{H}_T$, $T > t$, such that $\iota(s) \neq i'$ for $t < s < T$ and $\iota(T) = i'$. By Claim 11 there exists a finite such history which is reached with strictly positive probability where $\lambda_T \geq \hat{\lambda}(\mathcal{H}_t)$. Notice that $\hat{\lambda}(\mathcal{H}_T)$ depends on $\mathcal{H}_T$ only through θ_T^G, and given the way we have constructed $\mathcal{H}_T$: $\theta_T^G = \theta_t^G$. Thus, $\hat{\lambda}(\mathcal{H}_T) = \hat{\lambda}(\mathcal{H}_t)$, and hence $\lambda_T \geq \hat{\lambda}(\mathcal{H}_T)$.

∎

Proof of Theorem 2: If the expert is a good expert then the Law of Large Numbers yields $\lim_{t \to \infty} \theta_t^G(0) = \theta$, almost surely. Both parts of the theorem now follow from Claim 11.

∎

7.3 Proofs of Propositions 3 and 4 and Observations 2-5

Proof of Proposition 3: The threshold value $\hat{\lambda}$ is retrieved from (6)-(7), setting $\theta_t^G = 0$. Any value ε no less than the maximum reputational jump of an analyst with current reputation $\hat{\lambda}$ will satisfy the condition stated in the proposition.

∎

Proof of Observations 2-5: If $\sigma_t^0 < q_t^0$ then $\lambda_t^e(0, \underline{\sigma}_t, \underline{\xi}_t) > \lambda_t$, contradicting (8). Hence $\sigma_t^0 \geq q_t^0 = 1/2$, with strict inequality unless the market is trapped. This gives Observation 2.

Next, note that for any history $\mathcal{H}_t$ and in any equilibrium

$$\mathbb{P}_t(y_t \in \{a, b\}) = \frac{2(1-\mu)}{3} + \mu \sum_{(r,s)} \mathbb{P}_t(r, s) \big[\xi_t^a(r, s) + \xi_t^b(r, s)\big],$$

where by definition: $\xi_t^a(r, s) + \xi_t^b(r, s) \leq 1$, for all (r, s). When the market is trapped this last inequality becomes an equality. This gives Observation 3.

Observations 4 and 5 are immediate, since $p_t^0 = \mathbb{E}_t[\theta] = \lambda_t \theta_t^G$.

∎

Proof of Proposition 4: The arguments of the proof for break-down of screening follow along the lines of Theorem 1. Let $\beta_t(r, s)$ denote an investor's updated belief of τ from observing $(r_t, s_t) = (r, s)$. Observe that for all $r \in \{-1, 1\}$ and all t, $q_t^r \geq 1 - \phi > 0$. Thus, if $\pi < 1$, we can find $\epsilon > 0$ such that $|1 - \lambda_t| < \epsilon/2 \Rightarrow |1 - \beta_t(r, s)| < \epsilon$, for all $(r, s) \in \{-1, 1\} \times \{0, 1\}$. Letting ϵ tend to zero, $v_t(r, s)$ then approaches $\theta_t^G(r)$, for all $(r, s) \in \{-1, 1\} \times \{0, 1\}$. p_t^a (resp. p_t^b), by contrast, is bounded away from $\theta_t^G(1)$ (resp. $\theta_t^G(-1)$), so long as $\mu < 1$. These arguments establish that if $\pi < 1$ and $\mu < 1$ then screening necessarily breaks down, above a threshold level of reputation.

We next establish an analogue of Observation 1 from Section 3 and show that we can find $\epsilon(\cdot) : [0, 1) \to \mathbb{R}_+$, $\lim_{\delta \to 0} \epsilon(\delta) = 0$, such that if in equilibrium screening breaks down then $|\lambda_{t+1} - \lambda_t| \leq \epsilon(\delta)$. Fix δ, and consider an equilibrium in which screening breaks down. λ_{t+1} is then a deterministic function of r_t, which we may write $\lambda_{t+1}(r)$ for $r_t = r$. Moreover, $|\lambda_{t+1}(r) - \lambda_{t+1}(-r)|$ is strictly increasing in $\|\underline{q}_t - \underline{\sigma}_t\|$. It is now immediate to see from the form of the analyst's payoffs that fixing $\|\underline{q}_t - \underline{\sigma}_t\| > 0$ and letting δ tend to zero must either make payoffs from recommending $r_t = 1$ strictly greater than payoffs from recommending $r_t = -1$, or the converse. This is impossible, in equilibrium. Hence, for each $\delta \in [0, 1)$ we can find $\eta(\delta) > 0$, such that $\|\underline{q}_t - \underline{\sigma}_t\| < \eta(\delta)$, in any equilibrium. This in turn delivers, for each $\delta \in [0, 1)$, an $\epsilon(\delta) > 0$ such that $|\lambda_{t+1} - \lambda_t| \leq \epsilon(\delta)$, in any equilibrium. Furthermore $\lim_{\delta \to 0} \eta(\delta) = 0$ and so, evidently, $\lim_{\delta \to 0} \epsilon(\delta) = 0$.

$\blacksquare$

References

Anat R Admati and Paul Pfleiderer. A Monopolistic Market for Information. *Journal of Economic Theory*, 39(2):400–438, 1986.

Paul Asquith, Michael B Mikhail, and Andrea S Au. Information Content of Equity Analyst Reports. *Journal of Financial Economics*, 75(2):245–282, 2005.

Christopher Avery and Peter Zemsky. Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review*, 88(4):724–748, 1998.

Abhijit V Banerjee. A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 107 (3):797–817, 1992.

Roland Benabou and Guy Laroque. Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. *Quarterly Journal of Economics*, 107(3):921–958, 1992.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5): 992–1026, 1992.

Qi Chen, Jennifer Francis, and Wei Jiang. Investor Learning About Analyst Predictive Ability. *Journal of Accounting and Economics*, 39(1):3–24, 2005.

Martin W Cripps, George J Mailath, and Larry Samuelson. Imperfect Monitoring and Impermanent Reputations. *Econometrica*, 72(2):407–432, 2004.

Amil Dasgupta and Andrea Prat. Information Aggregation in Financial Markets with Career Concerns. *Journal of Economic Theory*, 143(1):83–113, 2008.

Jeffrey C Ely and Juuso Välimäki. Bad Reputation. *Quarterly Journal of Economics*, pages 785–814, 2003.

Lawrence R Glosten and Paul R Milgrom. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics*, 14(1): 71–100, 1985.

D Eric Hirst, Lisa Koonce, and Paul J Simko. Investor Reactions to Financial Analysts' Research Reports. *Journal of Accounting Research*, pages 335–351, 1995.

Ju Hu. Reputation in the Presence of Noisy Exogenous Learning. *Journal of Economic Theory*, 153:64–73, 2014.

In H Lee. Market Crashes and Informational Avalanches. *Review of Economic Studies*, 65(4): 741–759, 1998.

Xi Li. The Persistence of Relative Performance in Stock Recommendations of Sell-side Financial Analysts. *Journal of Accounting and Economics*, 40(1):129–152, 2005.

George J Mailath and Larry Samuelson. Reputations in Repeated Games. *The Handbook of Game Theory*, 4, 2013.

Roni Michaely and Kent L Womack. Market Efficiency and Biases in Brokerage Recommendations. *Advances in Behavioral Finance II*, pages 389–422, 2005.

Marco Ottaviani and Peter N Sørensen. The Strategy of Professional Forecasting. *Journal of Financial Economics*, 81(2):441–466, 2006.

Andreas Park and Hamid Sabourian. Herding and Contrarian Behavior in Financial Markets. *Econometrica*, 79(4):973–1026, 2011.

Canice Prendergast and Lars Stole. Impetuous Youngsters and Jaded Old-timers: Acquiring a Reputation for Learning. *Journal of Political Economy*, pages 1105–1134, 1996.

Sundaresh Ramnath, Steve Rock, and Philip Shane. The Financial Analyst Forecasting Literature: A Taxonomy with Suggestions for Further Research. *International Journal of Forecasting*, 24(1):34–75, 2008.

Jesper Rudiger and Adrien Vigier. Financial Experts, Asset Prices and Reputation. *Working Paper*, 2014.

Jesper Rudiger and Adrien Vigier. Asymmetric Information Acquisition in Financial Markets. *Mimeo*, 2015.

David S Scharfstein and Jeremy C Stein. Herd Behavior and Investment. *American Economic Review*, 80(3):465–79, 1990.

Lones Smith and Peter N Sørensen. Pathological Outcomes of Observational Learning. *Econometrica*, 68(2):371–398, 2000.

Ran Spiegler. The Market for Quacks. *Review of Economic Studies*, 73(4):1113–1131, 2006.

Brett Trueman. Analyst Forecasts and Herding Behavior. *Review of Financial Studies*, 7(1): 97–124, 1994.

Xavier Vives. *Information and Learning in Markets: The Impact of Market Microstructure.* China: Social Sciences Academic Press, 2012.

Thomas Wiseman. Reputation and Exogenous Private Learning. *Journal of Economic Theory*, 144(3):1352–1357, 2009.