Learning from Speculating and the No Trade Theorem

Dmitry Lubensky* Doug Smith †

October 1, 2017

Abstract

To overcome the no-trade theorems of Aumann (1976) and others, models of speculative trade have relied on agents that do not maximize expected wealth (noise traders). We develop an overlapping generations model in which rational wealth-maximizing speculators with a common prior trade a common value asset based only on private information. The rationale for trade is experimentation: an agent that trades learns about her type, exiting if it is low and continuing to benefit from future trades if it is high. We demonstrate that the learning motive always overcomes adverse selection, regardless of the rate of learning or the benefit of being a high type, generating a substantial volume of purely speculative trade. In a single period snapshot the no trade theorem would ensue, and observed trade is attributed to younger cohorts who appear as noise traders by entering unprofitable trades.

*Kelley School of Business, Indiana University. dlubensk@indiana.edu

†Federal Trade Commission. dsmith2@ftc.gov.

The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Trade Commission or any Commissioner.

1 Introduction

A classic result of economic theory is that asset trade cannot occur solely for informational reasons. This fact was initially formalized by Aumann (1976) and sparked a robust literature which continued to find this result in increasingly complex settings (e.g. Milgrom and Stokey (1982), Kreps (1977), and Tirole (1982)). As it turns out, regardless of how one models a market, rational agents cannot trade for purely speculative purposes. Tirole states that even in a dynamic setting, "speculation relies on inconsistent plans and is ruled out by rational expectations".

The driving force behind these results is adverse selection. A buyer of an asset is concerned that the seller trades only when she has some private negative information, and adjusts by trading only when her own signal is sufficiently positive. The seller does the same, and in making these adjustments the two traders exacerbate the selection problem. By Aumann's logic in equilibrium this eventually leads to full unravelling. Trade implies common knowledge of both the buyer and seller not losing money, which leads to a contradiction when trades are zero-sum.

To overcome adverse selection a speculator must thus believe that her counterparty systematically loses money. This can occur if the opposing trader has ulterior motives for asset ownership such as the hedging of risk or liquidity concerns (e.g. Kyle (1985), Glosten and Milgrom (1985)), or if the opposing trader is also a speculator but suffers from behavioral biases or a misunderstanding of the market environment (e.g. De Long et al. (1990)). Traders that lose money, often referred to in the literature as noise traders, provide the extra degree of freedom to generate trade, but the patterns of trade in noise trader models often do not match what we observe in financial markets. First, it is generally thought that observed volume is substantially higher than what can be attributed to classical motives such as risk-hedging (Odean (1999)), but adding privately informed speculators to a classical market should depress volume, not augment it. And second, noise traders must typically account for a substantial proportion of trading volume in order to lure in speculators wary of adverse selection.[1] Meanwhile, the overwhelming majority of transactions in financial markets involve sophisticated professionals, not overconfident day-traders or retirees cashing in pensions.[2] Thus the implication of noise trader models is that in

[1]For example, in Kyle's canonical single-period model informed traders account for exactly one half of the trading volume. Similarly, Easley et al. (2002) estimates that informed traders on average account for only 19% of volume for NYSE-listed equities. An exception is Mahani and Bernhardt (2007), to be discussed shortly, which finds for some parameters a large proportion of professional trading.

[2]A recent industry study by the TABB group (https://research.tabbgroup.com/report/v07-023-us-equity-high-frequency-trading-strategies-sizing-and-market-structure) finds that 88% of the equities trading volume on the NYSE came from high-frequency traders, institutional investors, and hedge funds, 11% from retail investors, and 1% from other sources. Furthermore, for many complex over-the-counter financial

financial markets a large proportion of professionals trade in a way that is inconsistent with profit maximization. Whether markets would force out such traders in the long run (Fama (1988)) or whether they may continue to survive (De Long et. al. (1990)) has been a topic of an ongoing debate.

In this paper we present an alternative approach and demonstrate that the presence of noise traders is not necessary to generate trade. Instead, we develop a model in which all agents are rational risk-neutral speculators that trade a common value asset based only on private information. The motive that helps agents overcome adverse selection is experimentation – by trading an agent not only receives a payoff but also learns about whether she is good at trading. Adverse selection is still present and agents sometimes enter trades that on average lose them money, however we show that the value of the information learned from the trade compensates for this loss.

In the model a cohort of new agents is born every period and each agent has the opportunity to transact in an existing pool. Agents have either a high or low type but do not directly observe this; instead each may learn about her type through trading. In every period agents receive private pre-trade signals about the value of an asset, with high types receiving signals that are more accurate. Agents are then randomly pairwise matched and choose whether to buy, sell, or abstain from trading. This interaction generates a stage payoff and provides each agent information about their type. For instance, an agent whose pre-trade signal predicts a high asset value but whose realized payoff reveals a low asset value may shift her posterior toward the low type. The period ends with a pool of traders, each holding a belief about their type that depends on their trading history. In the next period there is an inflow of a new cohort and an outflow of existing traders at an exogenous rate. We solve for the steady state equilibrium of this market.

The model accommodates several interpretations of an agent's type. For instance, an agent may have exclusive access to a source of information, such as an insider at a company whose stock is being traded, and the type is the quality of the source. Alternatively, there may be a very large number of public signals but due to cognitive or computational constraints agents may only implement strategies that condition on a smaller subset of the signals. Here an agent is an algorithm that bases trades on a subset of signals, and the agent's type reflects the informativeness of the subset. The type may also be person-specific and reflect an agent's skill, for instance the speed or accuracy at which the trader identifies and executes arbitrage opportunities regardless of the sources of information.

We consider the case in which assets have a purely common value, so that the sum of the payoffs of two matched traders is zero. The steady state pool is composed of agents with

products trading is legally restricted to only professionals.

varying beliefs about their type, and if agents in this pool played for one period only the no-trade theorem would arise due to the concern of trading against a counterparty with a higher type. However, in the dynamic environment even agents in the new cohort, who on average face a counterparty with a higher type and expect to lose money in the current period, engage in trade. Their rationale is that if through trading they learn that their type is high they can continue to trade in future periods and earn a positive continuation value, while if they learn that their type is low they can cut their losses by abstaining from future trades. We demonstrate that in the zero sum environment the information value of the first trade exactly offsets the losses from adverse selection and thus an equilibrium is supported in which all new agents as well agents from older cohorts trade.

The intuition for this result is that new agents are matched against future versions of themselves. That is, in steady state the probability that a new agent is matched with a counterparty with a higher type in her first period equals the probability that in a later period this agent has the higher type and her counterparty is that period's new agent. More broadly, in any pairwise match between traders with two different histories, in steady state a new agent is equally likely to be either counterparty over her lifetime, and since trades are zero sum her expected lifetime payoff is zero. New agents thus lose money on the first trade but make it back later on, and abstaining in the first period does not provide a profitable deviation.

We show the existence of purely speculative trade quite generally, and in particular the result does not depend whether one learns very little (including nothing) or a lot from trading or the magnitude of the stage payoff advantage of high types. In fact, the only truly necessary assumption is that agents cannot learn about their type without attempting to trade. We believe this assumption is quite natural, since in many professional financial markets the success of a trading strategy is not easily observed counterfactually. In over the counter markets for example, a potential buyer of an asset such as a credit default swap does not observe the price at which the asset would have traded without approaching a counterparty to underwrite and sell the product. Furthermore, even for exchange-traded assets the success of a trading algorithm cannot be observed without implementation, because prices and market depth would react to the trades the algorithm would initiate.[3]

In a single period snapshot of our dynamic model, new agents willingly lose money on that period's trade and thereby are observationally similar to the noise traders in the classic models such as Kyle (1985) and Glosten and Milgrom (1985). In this sense, one interpretation of our model is that it identifies learning about one's skill as a foundation

[3]Consider a spoofing algorithm, i.e. a sequence of buy orders designed to move up the market price followed by a sell order at the inflated price. Typically one cannot learn whether the market price is driven up sufficiently for the sell order or whether the prices are sufficiently low for the buy orders without actually using the algorithm.

for the existence of noise traders. Because over their lifetimes our noise traders do not lose money, we avoid the concern that markets ought to weed out ineffective traders. In fact, our model directly describes the way in which this selection occurs. Furthermore, we highlight the fact that it is not only speculators, i.e. in our model the experienced traders who on average make money on the current trade, but also the noise traders that make decisions based on information about the value of the asset. This is a useful distinction when evaluating how well the market aggregates available information, or how changes in the uncertainty over the asset's value affect trading patterns, or a myriad of similar lines of inquiry.

An important related work is Mahani and Bernhardt (2007), which also studies an overlapping generations model where speculators learn about their abilities by trading. Their model includes liquidity traders, while speculators must expend costs to gather information, and the liquidity traders' losses finance this activity.[4] Although trading volume is not the authors' main focus, they show numerically that for some parameter values in which the advantage of informed speculators over liquidity traders is most pronounced, speculative trade can comprise a large proportion of total volume. Thus, one interpretation is that speculative trade must either be supported by a large number of liquidity traders making small losses or a small number of liquidity traders making large losses. In the present paper we argue that neither is necessary, and that in fact speculative trade is self-supporting. A young speculator does not need to make money from liquidity traders to offset losses to more experienced speculators, because the information value of trading is sufficient.

Several recent empirical studies support the idea that learning about one's skill as a trader provides an important rationale for trading. Barber et al. (2010) demonstrates that traders expect future performance to be correlated with their past performance and are more likely to exit when initially losing money. Seru et al. (2010) finds a similar result at the individual level and also shows that older cohorts outperform younger cohorts, mostly due to the selection inherent in the exit decision. Some theoretical work has also examined the role of learning in speculative trade. Bond and Eraslan (2010) supports information-based bilateral trade by having the process of trading release information that improves the productive value of the asset for the owner. In our work the process of trade is also necessary to release information, but by contrast the information is not productive as there are no gains from trade in the market. Closer in spirit is Gervais and Odean (2001), in which speculators learn about their skill as they trade in multiple periods. The paper focuses on boundedly rational speculators that over-attribute prior success to skill, and demonstrate that overconfidence is strongest early in a speculator's career and diminishes over time. In our setting a similar dynamic occurs but due to the

[4]Indeed no trade can occur in this setting when there are no liquidity traders.

overlapping generations structure we can support trade without the use of noise traders and generate behavior that is rational in the long run but seemingly overconfident in the short run.

The rest of the paper proceeds as follows. Section 2 presents an illustrative example in which we describe a simple information structure such that purely speculative zero-sum trade is supported in steady state. In this example, the result that a new agent's first period loss from trading in an adversely selected pool is exactly offset by the information value emerges without relying on the explicit characterization of the equilibrium strategy or the steady state pool. To demonstrate that this is the case more broadly, we estbalish in Section 3 that in any steady state the lifetime payoff of an agent is zero when stage games are zero sum, describe the more general asset trading model in Section 4, and establish the existence of a steady state equilibrium with positive trade in Section 5. Section 6 then provides reasons why the equilibrium with trade may be more likely to occur than other equilibria, and Section 7 concludes.

2 Learning about Ability: An Example

In this section we study an overlapping generations setting in which agents live exactly two periods and in every period are matched in pairs to trade an asset of uncertain value. Each agent has either a high or low type, which corresponds to the quality of information she receives about the asset, but does not know her type when she is born. By trading the agent observes how the realized asset value compares to her pre-trade signal and draws inference from this about her type. We demonstrate (i) that there exists a steady state equilibrium with a substantial amount of trade and (ii) that due to adverse selection no trade would occur in the single-period snapshot of this environment.

The goal of the example is to capture many of the important features of the main model in a simpler setting and provide an outline of the arguments used for the general result. Some assumptions, for instance those regarding the distributions of asset values and the correlation of agents' pre-trade signals, are made purely to facilitate exposition and relaxed in the main section. Other assumptions, such as the fact that agents cannot learn about their type without attempting to trade, are central to the main result and are highlighted as such in the analysis.

Entry and exit. Time flows discretely in periods (..., t-1, t, t+1, ...). Each period a unit mass of new agents enter and each agent i independently draws a type

$$\theta_i = \begin{cases} G & \text{with probability } \frac{1}{2} \\ B & \text{with probability } \frac{1}{2} \end{cases}.$$

Types correspond to whether an agent is good (G) or bad (B) at obtaining or interpreting information about the value of an asset, as described shortly. Entering agents do not directly observe their type but may learn about it by interacting in the market. We denote by μ an agent's belief that she is type G, with all new agents starting with the prior $\mu_0 = \frac{1}{2}$.

Agents live exactly two periods. We refer to agents in their first period as young and to agents in their second period as old. The pool of agents is composed of these two overlapping generations and all agents in the pool are pairwise matched regardless of age.

Asset Value. In any given period, an asset has a value $v = v_1 + v_2$, in which v_1 and v_2 are independent stochastic components. Each component v_i is distributed on $[-1, 1]$ according to state-dependent density $\phi(v_i|\alpha_i) = \frac{1}{2}(1 + \alpha_i v_i)$, in which the high state $\alpha_i = 1$ and the low state $\alpha_i = -1$ are equally likely. The two densities are triangle distributions, downward sloping in the low state and upward sloping in the high state, and symmetric with respect to each other around $v_i = 0$, as depicted in Figure 1.

Stage game. In every period agents are randomly pairwise matched, each receiving a private binary pre-trade signal $s_i \in \{-1, 1\}$ about the state α_i of her component of an asset. The quality of the signal depends on the agent's type, specifically

$$Pr(s_i = \alpha_i) = \begin{cases} 1 & \text{if } \theta_i = G \\ \\ \frac{1}{2} & \text{if } \theta_i = B \end{cases},$$

so that for type G the signal perfectly reveals state α_i and for type B the signal is pure noise. Having observed their pre-trade signals each agent chooses an action $x_i \in$ {buy, sell, out} and trade occurs if and only if one agent chooses "buy" and the other "sell". When a trade occurs the buyer's payoff is $v_1 + v_2$, the seller's payoff is $-(v_1 + v_2)$, and both observe the realizations v_1 and v_2. When no trade occurs then both agents do not observe the realizations of the components and receive a payoff of zero.

A Trading Equilibrium

We now show that despite the presence of adverse selection in the steady state pool and the fact that trades are zero-sum, there is an equilibrium in which a substantial amount of trade takes place.

Proposition 1 *There exists a steady state equilibrium in which*

i. every participating agent follows her signal ("buy" if $s_i = 1$, "sell" if $s_i = -1$), *ii. all young agents participate,*

iii. there exists $\hat{v} \approx 0.064$ so that old agents participate iff in the previous period they traded and either $s_i = 1$ and $v_i \geq \hat{v}$ or $s_i = -1$ and $v_i \leq -\hat{v}$.

We check whether this is an equilibrium in two steps. First we determine the composition of the steady state pool of participating agents if all agents follow the above strategy. Then we check whether the above strategy is a best response when the pool is thus composed.

Composition of the steady state pool. In steady state there are four kinds of participants: young and old agents of types G and B. The numbers of participating good and bad young agents are $N_{1G} = N_{1B} = \frac{1}{2}$ since in the conjectured equilibrium all young agents participate. The number of participating old agents of type G is given by

$$\begin{aligned} N_{2G} &= N_{1G} \cdot \Pr(\textit{trade when young} \,|G) \cdot \Pr(\textit{favorable realization} \,|G) \\ &= \frac{1}{2} \cdot \left(\frac{1 + N_{2G} + N_{2B}}{2} \cdot \frac{1}{2} \right) \cdot \left(1 - \frac{1}{4}(1 + \hat{v})^2 \right), \end{aligned} \tag{1}$$

reflecting in the first term that such an agent was born type G, in the second term that she was matched with a willing participant with an opposing pre-trade signal, and in the third term that she observed a sufficiently favorable post-trade signal. The expression for the third term follows from the fact that agents of type G receive a pre-trade signal that is fully informative of the state. In particular, if a type G agent receives the "buy" pre-trade signal $s_i = 1$, then the true state is $\alpha_i = 1$ and $\Pr(\textit{favorable realization} \,|G) = 1 - \Phi(\hat{v}|\alpha_i = 1) = 1 - \frac{1}{4}(1 + \hat{v})^2$, and by symmetry the same expression holds for the "sell" pre-trade signal $s_i = -1$. Similarly, the number of participating old agents of type B is given by

$$\begin{aligned} N_{2B} &= N_{1B} \cdot \Pr(\textit{trade when young} \,|B) \cdot \Pr(\textit{favorable realization} \,|B) \\ &= \frac{1}{2} \cdot \left(\frac{1 + N_{2G} + N_{2B}}{2} \cdot \frac{1}{2} \right) \cdot \left(\frac{1}{2}(1 - \hat{v}) \right), \end{aligned} \tag{2}$$

the third term now accounting for the fact that agents of type B get favorable post-trade signals with a different probability. Namely, if a type B agent receives the "buy" pre-trade signal $s_i = 1$ then the true state is either $\alpha_i = 1$ or $\alpha_i = -1$ with equal chance and $\Pr(\textit{favorable realization} \,|B) = \frac{1}{2}(1 - \Phi(\hat{v}|\alpha_i = 1)) + \frac{1}{2}(1 - \Phi(\hat{v}|\alpha_i = -1)) = \frac{1}{2}(1 - \hat{v})$. Inspection of these expressions confirms that type G agents get favorable signals more often than type B agents, and thus there are more type G than type B old agents in the steady state pool.

The overall pool composition can be computed by solving the system of equations (1) and (2) in terms of $\hat{v}$ and then plugging in to obtain

$$\mu_s(\hat{v}) \equiv \frac{N_{1G} + N_{2G}}{N_{1G} + N_{2G} + N_{1B} + N_{2B}} = \frac{33 - \hat{v}^2}{64}, \tag{3}$$

in which μ_s is the proportion of participating agents of type G or equivalently the average belief of a participating agent. Observe that when $\hat{\nu} = -1$ every old agent that traded participates and when $\hat{\nu} = 1$ no old agents participate, and in either case the average belief in the pool is one half. For all interior thresholds the average belief among participants in the pool that one is a high type is strictly higher than one half and thus a new agent is on average of a lower type than her counterparty. We now verify that the conjectured equilibrium strategy, including participation by young agents, is a best response.

Stage payoffs. For an agent of type G the expected value of her component is $E[\nu_i|\alpha_i = 1] = \frac{1}{3}$ when she is a buyer and $E[\nu_i|\alpha_i = -1] = -\frac{1}{3}$ when she is a seller. For an agent of type B the expected value of her component equals the unconditional expected value $\frac{1}{2}E[\nu_i|\alpha_i = 1] + \frac{1}{2}E[\nu_i|\alpha_i = -1] = 0$ regardless of whether she is a buyer or seller, since her pre-trade signal is uninformative. Consequently, a buyer who has type G with probability μ facing a seller of type G with probability μ_s has a conditional expected trading payoff of

$$u_i(\mu, \mu_s) = E[\nu_i|\mu] + E[\nu_j|\mu_s] = \frac{1}{3}(\mu - \mu_s),$$

and by symmetry the same holds for a seller.

Beliefs. All young agents have the prior $\mu_0 = \frac{1}{2}$, while old agents form beliefs from their trading history using Bayes rule. An old agent that did not trade in the first period retains their prior of $\mu = \frac{1}{2}$.[5] Otherwise, an old agent that in the previous period received a buy signal $s_i = 1$, bought the asset, and observed post-trade signal ν_i has belief

$$\mu(\nu_i|s_i = 1) = \frac{\phi(\nu_i|\alpha_i = 1)}{\phi(\nu_i|\alpha_i = 1) \ + \ \frac{1}{2}\Big(\phi(\nu_i|\alpha_i = 1) + \phi(\nu_i|\alpha_i = -1)\Big)} \ = \ \frac{1+\nu_i}{2+\nu_i}. \tag{4}$$

As depicted in Figure 1, for a buyer higher realizations of ν_i result in a higher belief, with the posterior exceeding the prior of one half whenever $\nu_i \geq 0$. Any $\nu_i \geq 0$ is thus good news for the buyer, but as we will demonstrate not all good news is good enough to induce participation, i.e. $\hat{\nu} > 0$. A seller's belief is symmetric around zero relative to the buyer's, thus $\mu(\nu_i|s_i = -1) = \mu(-\nu_i|s_i = 1) = \frac{1-\nu_i}{2-\nu_i}$.

Strategies. First we verify that the proposed threshold strategy is a best response for old agents. Since old agents have no continuation they participate only if their expected stage payoff is positive, i.e. if their posterior is higher than the pool average. For buyers

[5] An agent that tries and fails to trade learns that her counterparty did not take the opposing position, which is potentially informative of the counterparty's pre-trade signal. In the current example agents' signals are independent, thus learning her counterparty's signal does not affect the agent's belief about her own type. In the main model signals are allowed to be correlated and the agent may update her beliefs even when no trade occurs.

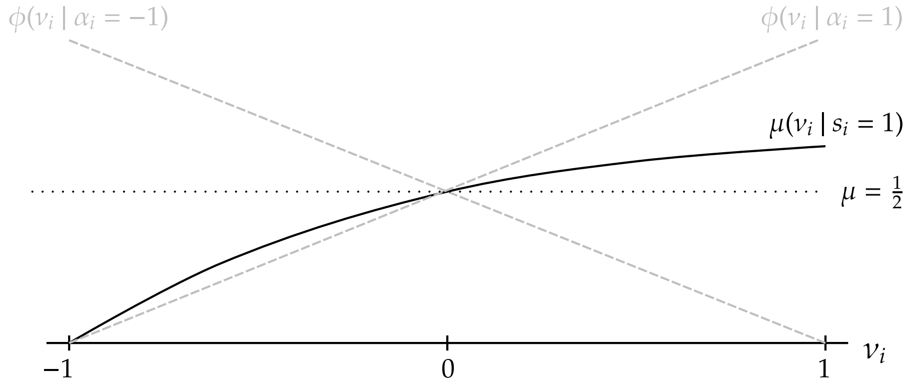

Figure 1: Belief conditional on buying

this implicitly defines the threshold post-trade signal so that $\mu(\hat{\nu}|s_i = 1) = \mu_s(\hat{\nu})$, which using (4) and (3) yields $\hat{\nu} \approx 0.064$ as a solution.[6] In other words, if old agents in the steady state pool use $\hat{\nu}$ as a threshold then it is also a best response for an old buyer to use $\hat{\nu}$ as a threshold when facing this pool. Finally observe that $\mu(\hat{\nu}|s_i = 1) = \mu(-\hat{\nu}|s_i = -1)$ and thus $-\hat{\nu}$ is also the threshold for sellers.

Note that $\mu_s(\hat{\nu}) \approx 0.52 > \mu_0$, implying that old agents must receive sufficiently positive news in the first round in order to participate. The threshold old agent is by definition at a disadvantage with respect to all other old agents, and in order for her payoff to be non-negative she must hold an advantage over new agents. By this argument, an old agent that did not trade does not expect to make money from new agents and expects to lose money to old agents, and therefore optimally abstains.

Having thus verified that the old traders' strategy is a best response, it remains now to demonstrate that it is a best response for all young agents to participate. First, we confirm that young agents that participate follow their signals (i.e. buy if $s_i = 1$ and sell if $s_i = -1$). The expected stage payoff from following one's signal is higher than from taking the opposite action, however playing against one's signal could potentially have a higher information value. However in the current example this is not the case. An agent's inference depends only on whether a trade occurs, not on whether she was a buyer or a seller in the trade. Since asset components are independent, choosing "buy" or "sell" induces a trade with the same probability, and thus both actions are equally informative. Thus there is a stage payoff advantage and no informational disadvantage to following one's signal.

[6]In fact there is an exact solution, with value $\hat{\nu} = \frac{1}{3}\left(\sqrt[3]{3\sqrt{88197} + 298} - \frac{89}{\sqrt[3]{3\sqrt{88197}+298}} - 2\right)$.

The lifetime payoff of a young agent that initially participates is given by

$$V = \Pr(trade)\Big(\Pr(G)\big[u_i(1,\mu_s) + \Pr(\nu_i s_i \geq \hat{\nu}|G)\Pr(trade)u_i(1,\mu_s)\big]$$
$$+ \Pr(B)\big[u_i(0,\mu_s) + \Pr(\nu_i s_i \geq \hat{\nu}|B)\Pr(trade)u_i(0,\mu_s)\big]\Big).$$

The young agent, if matched with a participating agent with an opposing signal, receives a stage payoff and conditional on receiving a favorable realization (i.e. if $\nu_i s_i \geq \hat{\nu}$) may trade again in the next period. We re-arrange the expression by factoring out the first term of the outer parentheses and obtain the following:

$$V = \Pr(trade)u_i(1,\mu_s)\Pr(G)\big(1 + \Pr(\nu_i s_i \geq \hat{\nu}|G)\Pr(trade)\big)$$
$$\times \left(1 + \frac{u_i(0,\mu_s)}{u_i(1,\mu_s)} \frac{\Pr(B) + \Pr(B)\Pr(trade)\Pr(\nu_i s_i \geq \hat{\nu}|B)}{\Pr(G) + \Pr(G)\Pr(trade)\Pr(\nu_i s_i \geq \hat{\nu}|G)}\right)$$

Using the fact that $u_i(\mu,\mu_s) = \frac{1}{3}(\mu - \mu_s)$ and the definitions of $N_{1\theta}$ and $N_{2\theta}$, the expression is further simplified:

$$= \Pr(trade)u_i(1,\mu_s)\Pr(G)\big(1 + \Pr(\nu_i s_i \geq \hat{\nu}|G)\Pr(trade)\big)$$
$$\times \left(1 - \frac{\mu_s}{1-\mu_s} \frac{N_{1B} + N_{2B}}{N_{1G} + N_{2G}}\right)$$

$$= \Pr(trade)u_i(1,\mu_s)\Pr(G)\big(1 + \Pr(\nu_i s_i \geq \hat{\nu}|G)\Pr(trade)\big)$$
$$\times \left(1 - \frac{\mu_s}{1-\mu_s} \frac{1-\mu_s}{\mu_s}\right)$$
$$= 0$$

This calculation shows that participating when young gives as high as payoff as abstaining, and is thus a best response. Therefore, the strategies of young and old agents are optimal and this concludes the proof.

Observe that demonstrating that a young agent is indifferent to participating in the adversely selected pool did not rely on the particular value of the threshold $\hat{\nu}$ or on the distributions of asset component values $\phi(\nu_i|\alpha_i)$. We will demonstrate in the main model that this is a general feature and applies in a broader setting, which enables us to prove the existence of a steady state with trading without explicitly describing either the equilibrium strategy or the steady state pool composition.

Proposition 2 *(No Trade Theorem) In every period of the overlapping generations model, approximately* 59% *of agents actively trade, while if all agents in the steady state pool played a one-shot game then* 0% *would actively trade.*

Proof In the previously described equilibrium, plugging $\hat{\nu} \approx 0.064$ into (1) and (2) yields that the total number of active agents in a given single period is

$$N(active) = N_{1G} + N_{1B} + N_{2G} + N_{2B} \approx 1.174,$$

and since the total mass of agents present equals two, the proportion of active traders is approximately 59%.

To show that no trade would occur in a single period snapshot, let $Z(\mu)$ denote the distribution of posteriors in the steady-state equilibrium, and observe that there is an atom at $\mu = \frac{1}{2}$ that corresponds to young agents and full support on $\mu \in \left[0, \frac{2}{3}\right]$ corresponding to old traders having received different post-trade signals. In any equilibrium of the one shot game, both buyers and sellers must follow threshold strategies, i.e. there exists a belief $\overline{\mu}_b$ for buyers and $\overline{\mu}_s$ for sellers so that agents agree to participate only if their belief is higher than their respective threshold. The marginal seller expects to trade with the average buyer and the marginal buyer expects to trade with the average seller, and both marginal types are willing to participate only if

$$E[\nu \, | \mu_b \geq \overline{\mu}_b, \mu_s = \overline{\mu}_s] \leq 0 \leq E[\nu \, | \mu_b = \overline{\mu}_b, \mu_s \geq \overline{\mu}_s].$$

However, since $E[\nu \mid \mu_b, \mu_s]$ increases in μ_b and decreases in μ_s, it must be that

$$E[\nu \, | \mu_b \geq \overline{\mu}_b, \mu_s = \overline{\mu}_s] \geq \; E[\nu \, | \mu_b = \overline{\mu}_b, \mu_s = \overline{\mu}_s] \; \geq \; E[\nu \mid \mu_b = \overline{\mu}_b, \mu_s \geq \overline{\mu}_s].$$

Both sets of inequalities hold only if

$$E[\nu \mid \mu_b \geq \overline{\mu}_b, \mu_s = \overline{\mu}_s] = E[\nu \, | \mu_b = \overline{\mu}_b, \mu_s \geq \overline{\mu}_s],$$

which in turn implies that $\overline{\mu}_b = \overline{\mu}_s = \frac{2}{3}$, the maximal value of the support of $Z(\mu)$. Thus trade can occur only between the maximally optimistic buyer and the maximally pessimistic seller, and the probability that two such traders are matched is zero. ■

Finally, observe in steady state approximately 85% of all active traders are young while only the remaining 15% are old. Recalling that young traders on average lose money in the stage while old traders make money, one may mistakenly infer from the steady state snapshot that 15% are sophisticated agents that trade based on information while the remainder are potentially uninformed noise traders. In fact, similar estimates of the percent of informed trade have been found in the literature (Easley et al. (2002)), and while this example is of course highly stylized, it is worth noting that our approach can generate seemingly uninformed trading in magnitudes close to those observed empirically.

The example presented in this section demonstrates that a substantial amount of trade can occur for purely informational purposes, as more than half of the agents actively trade despite there being no aggregate value from doing so. This may appear surprising in light of the no trade theorems, and our aim is now to explore which assumptions were necessary for it to obtain. As we will demonstrate, many of the particulars of the example, such as the functional form of the distributions of asset value, the independence of signals, the length of an agent's life, and even the connection between stage payoffs and inference about skill are irrelevant for the existence of trade. The key feature is that information about one's type may only be obtained by attempting to trade.

3 A Steady State Result

We first take a step back and consider a more general overlapping generations setting where in each period agents are randomly pairwise matched and the interaction is zero sum but arbitrary otherwise. In such a setting we demonstrate that if all agents use the same dynamic strategy σ, then when the pool is in steady state an agent's expected lifetime payoff exactly equals zero. Then, in the next section which models the stage game as that of two speculators trading a common value asset, this result is used to demonstrate that it is a best response for new agents to trade even though the steady state pool is adversely selected.

Time flows discretely in periods (..., t-1, t, t+1, ...). Each period a unit mass of ex-ante identical agents is born, who thereafter exit at an exogenous rate $\delta \in [0, 1]$ per period and live at most $T \leq \infty$ periods. Let M denote the expected lifetime and assume that it is finite, i.e. that $\min\left(\frac{1}{\delta}, T\right) < \infty$.

In every period agents in the pool are randomly pairwise matched and play a stage game, resulting in a payoff and a signal for each. Let history h_τ describe all the information observed by an agent in the first τ periods and let H be the set of all histories. An agent's strategy $\sigma : H \to \Delta(A)$ is a mapping from histories to probability distributions over the set of actions A available in any stage.

A pool is described by two objects: a probability measure $\rho(h)$ of agents at every history h and a strategy $\sigma(h)$ for each of those histories. A new agent that faces a fixed pool (σ, ρ) over her lifetime and plays the same strategy σ induces a measure over future histories $m(h|\sigma, \rho)$ and obtains an expected stage payoff $u(h, h'|\sigma, \rho)$ in every stage when her history is h and her counterparty's history is h'. Stage interactions are zero sum, so that $u(h, h'|\sigma, \rho) + u(h', h|\sigma, \rho) = 0$, and an agent's expected lifetime payoff is

$$V(\sigma, \rho) = \int_h \left(\int_{h'} u(h, h'|\sigma, \rho)\rho(h')\, dh' \right) m(h|\sigma, \rho)\, dh. \tag{5}$$

A steady state requires that the measure of agents at every history remains fixed over time. Since every history of length τ is populated only by agents that were born τ periods ago, the market is in a steady state if $m(h_\tau|\sigma, \rho) = M\, \rho(h_\tau)$ for every history h_τ. In other words, there is a steady state if the total number of agents in the pool with history h_τ equals the probability that a new agent reaches this history τ periods from now.

Proposition 3 *In any steady state* (σ, ρ)*, if stage games are zero sum then* $V(\sigma, \rho) = 0$*.*

Proof As written in (5), an agent's lifetime payoff is computed by integrating over the agent's histories on the outside and over her counterparty's histories on the inside. Alternatively, one may reverse the order as follows,

$$V(\sigma, \rho) = \int_h \left(\int_{h'} u(h', h|\sigma, \rho) m(h'|\sigma, \rho)\, dh' \right) \rho(h)\, dh, \tag{6}$$

integrating over own histories on the inside and the counterparty's histories on the outside. Adding (5) and (6) yields

$$\begin{aligned} 2V(\sigma, \rho) &= \int_h \left(\int_{h'} u(h, h'|\sigma, \rho) \Big(m(h|\sigma, \rho)\rho(h') - m(h'|\sigma, \rho)\rho(h) \Big)\, dh' \right) dh \\ &= \int_h \left(\int_{h'} u(h, h'|\sigma, \rho) \Big(M\rho(h)\rho(h') - M\rho(h')\rho(h) \Big)\, dh' \right) dh \\ &= 0, \end{aligned}$$

where the second line follows from the fact that games are zero sum (i.e. $u(h, h'|\sigma, \rho) = -u(h', h|\sigma, \rho)$) and that the market is in steady state (i.e. $m(z|\sigma, \rho) = M\, \rho(z)$ for all $z \in H$). ■

Another way to understand the preceding proposition is that for any two histories h and h', in steady state it is exactly as likely that the agent has history h and her counterparty has history h' as the reverse situation. Thus, if history h is more advantageous than history h', in steady state a new agent can expect to receive this advantage exactly as often as she expects to yield it over her lifetime. This is true in a steady state generated not only by a best response strategy σ^* but indeed by *any* strategy σ. Using this proposition, one does not need to explicitly characterize the optimal strategy as in the leading example to demonstrate that the lifetime payoff of a new agent is weakly positive.

We now return to the trading environment, explicitly describing the payoffs and information content of stage games, and use Proposition 3 to show (i) that there exists a steady state (σ^*, ρ^*) for which the strategy σ^* is a best response and (ii) that under σ^* all new agents trade.

4 Learning about Ability: A General Model

Entry and Exit. The dynamic structure is as described in the preceding section: time is discreet, a unit mass of agents is born each period, and each agent exits with probability δ in every period and lives at most T periods. Each agent i has a type

$$\theta_i = \begin{cases} G & \text{with probability} \quad \mu_0 \\ B & \text{with probability } 1-\mu_0 \end{cases},$$

where $\mu_0 \in (0,1)$. Entering agents do not observe their type directly, although they may learn about it by interacting in the market.

Stage Game. In every period each agent i receives a private binary pre-trade signal $s_i \in \{-1, 1\}$ and then chooses an action $x_i \in X \equiv \{\text{buy, sell, out}\}$. Agents are randomly pairwise matched and within each pair the function $q(x_1, x_2)$ determines whether trade occurs, with

$$q(x_1, x_2) = \begin{cases} 1 & \text{if } (x_1, x_2) = (\mathit{buy}, \mathit{sell}) \text{ or } (\mathit{sell}, \mathit{buy}) \\ 0 & \text{otherwise} \end{cases}.$$

When trade occurs the buyer's payoff is v and the seller's payoff is $-v$, when no trade occurs both agents receive a payoff of zero. At the end of the stage each agent i observes q, and conditional on trade also observes a post-trade signal $\gamma_i \in \Gamma$, which may capture payoffs and potentially other information related to the agents' types, as described in more detail below. To avoid redundancy we assume that agents do not directly observe stage payoffs. At the end of stage t an agent has thus observed $y_t \equiv (s_t, x_t, q_t, \gamma_t)$: her pre-trade signal s_t, her action x_t, whether trade occurred q_t, and if so her post-trade signal γ_t.

Asset value and signals. For each matched pair of agents of types θ_i and θ_j, the asset value v, pre-trade signals s_i and s_j, and post-trade signals γ_i and γ_j are drawn from a joint distribution $F(v, s_i, s_j, \gamma_i, \gamma_j | \theta_i, \theta_j)$ with the properties below:

(i) expected asset value equals zero:

$$E[v|\theta_i, \theta_j] = 0,$$

(ii) higher pre-trade signals correspond to higher asset value:

$$E[v|s_i = -1, s_j, \theta_i, \theta_j] < E[v|s_i = 1, s_j, \theta_i, \theta_j],$$

(iii) good agents get more meaningful pre-trade signals:

$$E[v|s_i = 1, s_j, G, \theta_j] - E[v|s_i = -1, s_j, G, \theta_j] > E[v|s_i = 1, s_j, B, \theta_j] - E[v|s_i = -1, s_j, B, \theta_j],$$

(iv) pre-trade signals symmetric in payoffs:

$$E[v|s_i, s_j, \theta_i, \theta_j] = -E[v| - s_i, -s_j, \theta_i, \theta_j],$$

(v) pre-trade signals symmetric in information:

$$\Pr(s_i, s_j, \gamma_i, \gamma_j \mid \theta_i, \theta_j) = \Pr(-s_i, -s_j, \gamma_i, \gamma_j \mid \theta_i, \theta_j).$$

Condition (i) is simply a normalization so that trade occurs at the ex-ante expected value of the asset. Condition (ii) is a definition of high and low pre-trade signals, while condition (iii) is similarly a definition of good and bad agent types. Observe that the leading example in which the bad agents' signal is a garbling of the good agents' signal constitutes a special case. Conditions (iv) and (v) ensure that a decision faced by an agent with a positive pre-trade signal looks identical to that faced by an agent with a negative pre-trade signal. As we will argue, conditions (iv) and (v) are not necessary for the existence of trade but rather simplify the presentation by putting agents with buy and sell signals on equal footing. Finally, the term $\Pr(\cdot)$ in condition (v) refers either to a discrete probability or a probability density, depending on the structure of the set of post-trade signals Γ.

In a given stage an agent can learn about her type in two ways. First, the mere fact of a trade occurring or not occurring can be informative. For example, if pre-trade signals s_i and s_j are correlated for good agents then if agent i attempts to trade but does not succeed, she may infer that it is likely both she and her counterparty received the same pre-trade signal, and thus both are likely good types. Second, when a trade does occur the agent also observes the post-trade signal γ_i, which in conjunction with her pre-trade signal s_i and her beliefs about her counterparty's pre-trade signal s_j provides further inference about her type θ_i using the joint distribution F.

The information structure is such that an agent that chooses x_i = "*out*" learns nothing in the stage. This is because she does not receive a post-trade signal and, by shutting down the possibility of trade, can draw no inference about her counterparty's signal s_j from the fact that trade does not occur. This plays an important role for our main result, since otherwise new agents might wait to learn for free rather than trade and incur a negative stage payoff to obtain the information. It may appear that such an assumption is restrictive, for instance in order to determine whether a particular trading strategy is effective one may abstain and simply observe market prices. However, in practice the counterfactual success of a trading strategy is seldom observed. Even with exchange-traded assets with public prices the trader cannot know how much her order would have moved the market to a different price, or alternatively whether her order would have been filled quickly enough before the price moved. Furthermore, in many over the counter markets it may be that no transaction takes place if the trader does not participate, thus she does not know if her order would have been accepted.

Strategies. Let $h = (y_1, ..., y_\tau)$ be an agent's history after τ stages and H the set of all possible histories. An agent's strategy is a mapping from histories and pre-trade signals to probability distributions over buying, selling, or abstaining:

$$\sigma : H \times \{-1, 1\} \to \Delta(X).$$

Because we impose symmetry between a buy and a sell signal in the distribution F, we can redefine the set of stage actions $A = \{-1, 0, 1\}$ as trading against one's signal, abstaining, or trading with one's signal. Then an agent's strategy simply maps from histories into mixed strategies over these re-defined actions,

$$\sigma : H \to \Delta(A),$$

and we let $\sigma(a, h)$ denote the probability that strategy σ assigns to pure action a for history h. Given the relabeling of actions we must also redefine the trading function as

$$q(a_i, s_i, a_j, s_j) = \begin{cases} 1 & \text{if } a_i s_i = -a_j s_j \text{ and } a_i, a_j \neq 0 \\ 0 & \text{otherwise} \end{cases}.$$

Stage Payoffs. An agent receives a nonzero stage payoff only when trade occurs, which in turn depends on the action of her counterparty. For instance, if both agents follow their pre-trade signals then trade occurs only if they receive opposing signals, and more generally trade occurs only if $s_j = -a_i a_j s_i$ (this also accounts for no trade occurring if either plays $a = 0$). It then follows that when an agent i with pre-trade signal s_i and type θ_i takes stage action a_i, her payoff when matched with a counterparty with type θ_j and action a_j is

$$w_i(a_i, \theta_i, a_j, \theta_j) = \Pr(s_j = -a_i a_j s_i | s_i, \theta_i, \theta_j) \cdot (a_i s_i) \cdot E[v | s_i, -a_i a_j s_i, \theta_i, \theta_j].$$

The first term uses the joint distribution F to compute the conditional probability that s_j is such that trade occurs, the second term determines whether i is a buyer or a seller, and the third term uses F to compute the conditional expected value of the asset. A trader faces an existing pool described by

$$P \equiv \Big\{P(a, \theta) \mid a \in \{-1, 0, 1\}, \theta \in \{B, G\}\Big\},$$

with $\sum_{a,\theta} P(a, \theta) = 1$ and each $P(a, \theta)$ denoting the proportion of traders that have type θ and take action a. Facing pool P, if agent i believes her type is G with probability μ then her expected payoff from action a_i is

$$u_i(a_i, \mu | P) \equiv \sum_{a_j, \theta_j} P(a_j, \theta_j)\Big(\mu w_i(a_i, G, a_j, \theta_j) + (1 - \mu) w_i(a_i, B, a_j, \theta_j)\Big). \tag{7}$$

Beliefs. While in the opening example an agent's stage inference about her type depended only on the relationship between her own pre- and post-trade signals, in principle the inference can also depend on her beliefs about the type and strategy of her counterparty. The likelihood of a stage observation $y = (s_i, a_i, q, \gamma_i)$ for an agent of type θ_i is

$$\ell(y \mid \theta_i, P) = \sum_{a_j, \theta_j} P(a_j, \theta_j) \left(\sum_{s_j} \mathbb{I}\Big(q(a_i, s_i, a_j, s_j) = q\Big) \Pr(s_i, s_j, \gamma_i | \theta_i, \theta_j) \right), \tag{8}$$

with corresponding likelihood ratio $\lambda(y|P) \equiv \frac{\ell(y|G,P)}{\ell(y|B,P)}$. The outer sum of the likelihood accounts for the probabilities of being matched with each of the six (a_j, θ_j) varieties of counterparty. The inner sum takes as given the counterparty's type and strategy and, for each of their pre-trade signals s_j which would have led to the observed trading outcome q, computes the joint likelihood of the pre- and post-trade signals (s_i, s_j, γ_i). For example, when agents i and j both follow their signals and a trade occurs and generates post-trade signal γ_i, it must be that $s_j = -s_i$ and the likelihood of this event (conditional on agents' types) is $\Pr(s_i, -s_i, \gamma_i | \theta_i, \theta_j)$. Meanwhile, in the same situation if agent j's strategy was to play "out" and no trade were observed, then both pre-trade signals s_j are consistent and the likelihood of this event is $\Pr(s_i, s_i, \gamma_i | \theta_i, \theta_j) + \Pr(s_i, -s_i, \gamma_i | \theta_i, \theta_j)$.

The likelihood ratio for a history $h = (y_1, ..., y_\tau)$ is simply the product of all the stage likelihood ratios and the prior likelihood ratio

$$\Lambda(h|P) = \frac{\mu_0}{1 - \mu_0} \prod_{t=1}^{\tau} \lambda(y_t|P),$$

with corresponding belief $\mu(h|P) = \frac{\Lambda(h|P)}{1+\Lambda(h|P)}$.

Dynamic Payoffs. We assume that agents maximize the expected value of lifetime income. Furthermore, we express asset values in net present value terms and impose that assets in the future have the same present value as assets today.[7] We wish to express an agent's lifetime payoff from using a strategy σ, and to this end we must first characterize the distribution over histories $m(h|\sigma, P)$ that her strategy σ induces. For a given history $h = (y_1, ..., y_\tau)$ and $t \leq \tau$, let $h_t = (y_1, ..., y_t)$ denote the first t observations in that history. A new agent that faces a pool P and uses strategy σ has a likelihood of reaching a history

[7] An alternative interpretation is that asset values grow over time, but that the rate of growth is exactly at the interest rate.

$h = (y_1, ..., y_\tau)$ given by

$$m(h|\sigma, P) = (1-\delta)^{\tau-1}\left(\mu_0 \prod_{t=1}^{\tau}\Big(\sigma(a_t, h_{t-1})\ \ell(y_t|G, P)\Big) \ +\ (1-\mu_0)\prod_{t=1}^{\tau}\Big(\sigma(a_t, h_{t-1})\ \ell(y_t|B, P)\Big)\right). \tag{9}$$

That is, conditional on the agent's type θ_i each observation $y_t \in h$ is an independent event, with the probability of action a_t prescribed by her strategy σ and the likelihood $\ell(y_t|\theta_i, P)$ derived as above from the primitive distribution F. Once again, $m(\cdot)$ is either a discrete probability or a probability density depending on the structure of the set of post-trade signals Γ.

An agent that faces pool P and follows strategy σ induces a distribution $m(h|\sigma, P)$ over future histories and an expected stage payoff $u_i\Big(\sigma(h), \mu(h|P) \mid P\Big)$ for each history, in which her inferred expected type at each history $\mu(h|P)$ depends on the pool P. Consequently, her expected lifetime payoff is

$$V(\sigma|P) = \int_{h\in H} u_i\Big(\sigma(h), \mu(h|P) \mid P\Big)\, m(h|\sigma, P)\, dh.$$

Steady State Equilibrium. A steady state requires that for every history the inflow and outflow of agents is identical. In particular, each agent today with a given history $\tilde{h}_\tau = (\tilde{y}_1, ..., \tilde{y}_{\tau-1}, \tilde{y}_\tau)$ will tomorrow transition to different history of length $\tau + 1$, and must be replaced by an agent with a history today of $\tilde{h}_{\tau-1} = (\tilde{y}_1, ..., \tilde{y}_{\tau-1})$. For there to be a steady state, it must be that the number of agents at $h_{\tau-1}$ and h_τ and the transition probability $\Pr(y_\tau|h_{\tau-1})$ from the former to the latter are such that the number of agents with history h_τ remains fixed.

An alternative way to understand steady state in our context is to observe that every history h_τ is populated exclusively by agents that were born τ periods ago, who reach this history with probability $m(h_\tau|\sigma, P)$ if over their lifetime they face a constant pool P. In the next period this history will be populated by agents born $\tau - 1$ periods ago, and with the same probability as long as those agents have also faced constant pool P. In other words, an equivalent definition of steady state is that P remains fixed over time. Letting $M = \frac{1-(1-\delta)^T}{\delta}$ denote the length of an agent's expected lifetime[8], we define steady state as follows.

Definition 1 *A strategy σ generates P if for each action $a \in \{-1, 0, 1\}$*

$$P(a, G) = \frac{1}{M}\int_{h\in H} \mu(h|P)\, \sigma(a, h)\, m(h|\sigma, P)\, dh$$

[8]Observe that $\lim_{\delta\to 0} M = T$ and $\lim_{T\to\infty} M = \frac{1}{\delta}$, thus if either $\delta = 0$ or $T = \infty$ let M be defined as one of these limits.

and

$$P(a, B) = \frac{1}{M} \int_{h \in H} (1 - \mu(h|P))\, \sigma(a, h)\, m(h|\sigma, P)\, dh.$$

By this definition, a strategy σ generates P if the proportion of agents taking action a with type θ equals the lifetime probability of a new agent having type θ and playing action a.

Definition 2 *(σ, P) is a steady state equilibrium if*

(i) σ maximizes $V(\sigma|P)$ and

(ii) σ generates P.

5 Equilibrium Existence and Properties

We now proceed to the paper's main result, using the fact from Proposition 3 that an agent's lifetime payoff in any steady state is non-negative to show that it is a best response for new agents to participate. To do this, we consider a strategy σ^* in which agents are restricted to participate in their first period and best respond thereafter and a pool P^* that is generated by σ^*, and argue that (σ^*, P^*) exists. Then we demonstrate that given P^* an agent that abstains in the first period does not increase her payoff – instead she receives a stage payoff of zero and a continuation payoff of at most zero. Thus (σ^*, P^*) is shown to be a steady state equilibrium and in this equilibrium trade occurs with positive probability.

Let $\Sigma_+ \equiv \{\sigma \mid \sigma(a = 0, h_0) = 0\}$ be the set of all strategies in which a new agent (history h_0) does not abstain.

Lemma 1 *There exists (σ^*, P^*) in which σ^* generates P^* and $\sigma^* \in \arg\max_{\sigma \in \Sigma_+} V(\sigma|P^*)$.*

The proof, found in the Appendix, relies on Kakutani's fixed point theorem. For a given pool P an agent has a set of optimal strategies $BR_+(P) \equiv \{\sigma \mid \sigma \in \arg\max_{\sigma \in \Sigma_+} V(\sigma|P)\}$, and each element $\sigma \in BR_+(P)$ then generates another steady state pool Q. Letting $\mathcal{P}$ be the set of possible pools, this process defines the correspondence $T : \mathcal{P} \to \mathcal{P}$. The set $\mathcal{P}$ is a probability simplex over the six (a, θ) combinations in the pool, and thus it is compact and convex. We then show that the mapping T is upperhemicontinuous, convex-valued, closed, and non-empty. This allows us to apply Kakutani's theorem and demonstrate that T has a fixed point.

Lemma 2 *Given a pool P^* generated by strategy $\sigma^* \in \Sigma_+$, if σ^* is optimal among strategies in which a new agent participates then σ^* is optimal among all strategies.*

Proof An agent that faces a steady state P^* solves a dynamic optimization problem in which her belief μ and her maximum remaining number of periods τ are sufficient state variables. For such an agent let $U(\mu, \tau|\sigma, P^*)$ denote the payoff of strategy σ in state (μ, τ), and let $U^*(\mu, \tau|P^*)$ denote the maximal payoff. Consider a conjectured optimal deviation strategy σ' in which a new agent abstains with positive probability $\sigma'(a = 0, h_0) > 0$. Then,

$$U^*(\mu_0, T|P^*) \equiv U(\mu_0, T|\sigma', P^*) = 0 + (1 - \delta)U^*(\mu_0, T - 1|P^*) \leq (1 - \delta)U^*(\mu_0, T|P^*).$$

Above, the equality follows because by assumption the optimal strategy σ' places positive weight on initially abstaining, and thus placing all the weight on initially abstaining is also optimal. This yields a stage payoff of zero and a continuation with the belief μ_0 and the maximal number of remaining periods $T - 1$. The ensuing inequality then stems from the fact that $U^*(\mu, \tau|P^*)$ increases in τ, since an agent with $\tau' > \tau$ remaining periods may simply mimic the strategy of having τ periods remaining and abstain in the final $\tau' - \tau$ periods. The implied inequality between the first and last term then yields $U^*(\mu_0, T|P^*) \leq 0$. But by Proposition 3, the payoff to following the strategy σ^* which generates the pool P^* gives a payoff of exactly zero, thus σ^* is optimal. ■

The following proposition combines the results of the previous two lemmas.

Proposition 4 *There exists a steady state equilibrium with strictly positive trading volume in which all new agents participate and all sufficiently confident experienced agents participate.*

Proof That a steady state equilibrium exists in which all new agents participate (and thus trading volume is strictly positive) is established by Lemmas 1 and 2. To see that experienced agents use a threshold strategy, observe that for any history with τ periods remaining if an agent with belief μ_1 participates then her continuation value is weakly positive. Then an agent with belief $\mu_2 > \mu_1$ can mimic the agent with belief μ_1 and obtain a higher payoff, and thus she also participates. ■

The stationary equilibrium and the no trade theorem

We note that the no trade theorem would obtain in a single period snapshot of our steady state equilibrium. That is, fix a steady state equilibrium (σ^*, P^*) in which all new agents participate, and the associated distribution over histories $\rho(h)$ and beliefs $\mu(h|P^*)$. In turn this implies a cumulative distribution over beliefs

$$Z(\mu) = \frac{1}{M} \int_{h \in H} \mathbb{I}(\mu(h|P^*) \leq \mu)\, \rho(h)\, dh.$$

Now imagine a single-period snapshot of this environment in which all agents draw a belief from $Z(\mu)$, play the stage game, and then realize the stage payoffs without continuation. In equilibrium, trading against one's own signal is dominated. Thus as in the leading example we look for an equilibrium in threshold strategies, that is a belief $\overline{\mu}$ so that all agents with a lower belief abstain and all agents with a higher belief follow their signal. Define $\mu_s \equiv E[\mu|\mu \geq \overline{\mu}]$ and observe that when trading with a random agent from the pool, the expected payoff of a buyer with threshold belief $\overline{\mu}$ is

$$E_{\bar{\mu},\mu_s}[\nu \mid \theta_i, \theta_j, s_i = 1, s_j = -1].$$

Since agents are anonymous and the expected payoff is zero whenever the two agents' types are the same, agent i's payoff can be rewritten as

$$\overline{\mu}(1-\mu_s)E[\nu \mid G, B, s_i = 1, s_j = -1] + (1-\overline{\mu})\mu_s E[\nu \mid B, G, s_i = 1, s_j = -1].$$

Finally, since the advantage of an agent of type G when matched with an agent of type B is independent of who is the buyer and who is the seller, the payoff further simplifies to

$$(\overline{\mu} - \mu_s)E[\nu \mid G, B, s_i = 1, s_j = -1].$$

It follows then that agent i's utility is non-negative only if $\overline{\mu} \geq \mu_s = E[\mu \mid \mu \geq \overline{\mu}]$. By definition this occurs only if $Z(\mu_s) = 1$, that is when the threshold belief is the maximal belief. This is the statement of the no-trade theorem. Observe finally that if post-trade signals γ are continuous (as in the leading example), then the probability of trade in the snapshot is exactly zero, while the probability of trade in the dynamic game is strictly positive since at least all new agents participate.

6 Equilibrium Selection

The previous section demonstrates the existence of an equilibrium where all new agents enter and continue to actively trade as long as their belief remains above threshold levels. However, initial entry with any probability is an equilibrium, including a probability of zero so that no trade takes place. Multiple equilibria are not uncommon of course, but given that initial entry is a weak best response, one may wonder what reason there is to expect the equilibrium above versus an equilibrium with no trade. To this end we briefly discuss two circumstances under which the trade equilibrium might be selected instead of the no-trade equilibrium.

First, consider an environment with a small number of agents that trade assets for classical reasons such as liquidity or risk-hedging concerns, and a large pool of speculators that are potential entrants. As long as the classical agents make up a positive proportion of the market, every speculator has a strict incentive to enter since they will on average capture some of the gains from trade. As entry persists and the proportion of speculators grows, the gains from trade per speculator diminish toward zero, and in the limit the environment is as the one described in the main model. Thus, while there is a continuum of equilibria at zero gains from trade, the limit of equilibria as gains from trade approach zero is the equilibrium with certain initial entry.

An alternative way to select the trading equilibrium is to examine the path to the steady state. Namely, consider a brand new market in which the first cohort of new agents makes up the entire pool. In this initial period the pool is not adversely selected so that in expectation the stage payoff from the first trade is zero and, due to the anticipated entry of new agents in the following period, the information value is strictly positive. Entry is thus a strict best response for this first cohort and, as experienced traders accumulate and the pool becomes more adversely selected, the payoff to entry remains positive but approaches zero as the pool approaches the steady state mix. Along the way toward steady state, not only do older generations make up a smaller proportion of the pool but also marginally skilled (in expectation) agents who otherwise would have exited in steady state remain while the pool is less selected, thus further slowing the adjustment.

By either of the previous arguments full trading can thus emerge as the unique equilibrium. This is not to say that full trading is the only reasonable equilibrium in out setting – certainly one may find other criteria which rule out full trading and instead select no trading. Rather, the aim is simply to demonstrate that the full trading equilibrium is likely not a pathological knife-edge outcome.

7 Conclusion

Classic no trade theorems imply that a rational speculator is only willing to trade based on her private information if she believes that her counterparty is willing to lose money. This observation has led to a class of models, starting with Kyle (1985) and Glosten and Milgrom (1985), in which a proportion of market participants known as noise traders either have preferences that are different from the speculators' and willingly lose money over their lifetime or alternatively misunderstand the market and either lose money or accept excessive risk. In this paper, we demonstrate that the presence of noise traders is not necessary to support information-based trade.

In our model all agents have the same preference for the asset and trade based on private information. The adverse selection of the no trade theorems is mitigated by agents' incentive to learn about their ability through participating in the market. An agent who learns that her ability is high remains in the market and continues to make profitable trades, while an agent who learns that her ability is low stops trading. The option value inherent in this information induces young agents to engage in trades with more experienced counterparties, trades on which young agents lose money on average. The learning motive thus not only rationalizes trade by young agents, but also alleviates adverse selection concerns for older agents who expect to transact with these younger cohorts. The key contrast from standard noise trader models is that while some agents are willing to lose money on a given trade, they are not willing to lose money over their lifetime. Thus, the concern that such traders would be selected out of the market due to their persistent losses does not apply here.

Some parallels exist between our framework and a model of learning by doing. It is straightforward to construct a learning by doing model where older agents that have participated in trade are more skilled than their younger counterparts, and younger agents accept short term expected losses from trade because they anticipate positive expected payoffs when they are old. The mechanism at work is different from our model, where an individual agent does not improve in ability over time, but agents are able to self select as they receive signals about their types. An important point of departure is that our model of learning *about* one's type draws a link between an agent's prior success and future willingness to trade, while no such link is immediate in a model of learning by doing.

Our model can also help explain the large volume of trade observed in securities markets. As noted in Odean (1999), "Trading volume on the world's markets seems high, perhaps higher than can be explained by models of rational markets" that attend to "investors' rebalancing and hedging needs" (p. 1279). Indeed, with the proliferation of high-frequency trading algorithms this point resonates even more; the vast majority of trading is done not by private investors but rather by professional speculators. But the presence of privately informed speculators should depress trading, not increase it, thus what accounts for the volume we observe?

One answer provided in the literature is that speculators can add to volume if they suffer from persistent behavioral biases, and thus participate in trades that are not in their interest (e.g. Odean (1999), De Long et al. (1990)). We provide an alternative explanation with rational experimentation as a motive. Professional speculators constantly seek trading opportunities, having to implement trading strategies in order to ascertain their efficacy and then exploiting the opportunities that they discover. We show that such constant experimentation is sustainable, and in particular that trying out a new strategy that will likely fail is a rational decision in a steady-state equilibrium.

References

Aumann, Robert J., "Agreeing to Disagree," *The Annals of Statistics*, 1976, *4* (6), pp. 1236–1239.

Barber, Brad, Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, "Do Day Traders Rationally Learn About Their Ability?," *Working Paper*, 2010.

Bond, Philip and Hlya Eraslan, "Information-based trade," *Journal of Economic Theory*, 2010, *145* (5), 1675 – 1703.

Easley, David, Soeren Hvidkjaer, and Maureen O'Hara, "Is Information Risk a Determinant of Asset Returns?," *The Journal of Finance*, 2002, *57* (5), 2185–2221.

Gervais, Simon and Terrance Odean, "Learning to be overconfident," *Review of Financial studies*, 2001, *14* (1), 1–27.

Glosten, Lawrence R and Paul R Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of financial economics*, 1985, *14* (1), 71–100.

Kreps, David M., "A note on fulfilled expectations; equilibria," *Journal of Economic Theory*, February 1977, *14* (1), 32–43.

Kyle, Albert S., "Continuous Auctions and Insider Trading," *Econometrica*, 1985, *53* (6), pp. 1315–1335.

Long, J Bradford De, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann, "Noise trader risk in financial markets," *Journal of Political Economy*, 1990, pp. 703–738.

Mahani, Reza and Dan Bernhardt, "Financial Speculators' Underperformance: Learning, Self-Selection, and Endogenous Liquidity," *The Journal of Finance*, 2007, *62* (3), 1313–1340.

Milgrom, Paul and Nancy Stokey, "Information, trade and common knowledge," *Journal of Economic Theory*, 1982, *26* (1), 17–27.

Odean, Terrance, "Do Investors Trade Too Much?," *American Economic Review*, 1999, *89* (5), 1279–1298.

Seru, Amit, Tyler Shumway, and Noah Stoffman, "Learning by trading," *Review of Financial studies*, 2010, *23* (2), 705–739.

Tirole, Jean, "On the Possibility of Speculation under Rational Expectations," *Econometrica*, 1982, *50* (5), pp. 1163–1181.

Appendix

Proof of Lemma 1: Existence of a steady state equilibrium with first period participation.

Let $\mathcal{P} \equiv \left\{P(a,\theta) \mid a \in \{-1,0,1\}, \theta \in \{G,B\}, \sum_{a,\theta} P(a,\theta) = 1, P(a,\theta) \geq 0\right\}$ be the probability simplex describing all possible pool distributions over actions a and types θ. Let Σ_+ be the set of strategies in which an agent initially participates, and let $BR_+(P) \equiv \{\sigma \mid \sigma \in \arg\max_{\sigma \in \Sigma_+} V(\sigma|P)\}$ be the collection of strategies that maximize an agent's payoff in this set. Then, define

$$T : \mathcal{P} \rightarrow \mathcal{P}$$

so that $Q \in T(P)$ if and only if there is a strategy $\sigma \in BR_+(P)$ so that for each $a \in \{-1,0,1\}$

$$Q(a,G) = \int_{h \in H} \mu(h|P)\ \sigma(a,h)\ m(h|\sigma,P)\ dh$$

and

$$Q(a,B) = \int_{h \in H} (1 - \mu(h|P))\ \sigma(a,h)\ m(h|\sigma,P)\ dh.$$

In other words, $Q \in T(P)$ if an agent facing a pool P and using a best response strategy $\sigma(P)$ generates a distribution over her future histories summarized by Q. The aim is to show that $T(P)$ has a fixed point, for which we rely on Kakutani's fixed point theorem.

First observe that $\mathcal{P}$ is a non-empty, convex, and compact subset of $\mathbb{R}^5$, thus satisfying the sufficient conditions for the domain in Kakutani's theorem. The key then is to focus on the continuity properties of T, in particular to demonstrate that T is non-empty, convex-valued, and closed (Claim 3), and that due to the fact that the value function is continuous in P (Claim 4), the mapping T is upperhemicontinuous (Claim 5).

Claim 3 $T(P)$ *is non-empty, convex-valued, and closed.*

Proof of Claim: That $T(P)$ is non-empty follows from the fact that for any pool P and any history h the set of available pure actions $A = \{-1,0,1\}$ is finite and at each history at least one of these pure actions is optimal. In turn implies the existence of the pool $Q \in T(P)$ generated by taking these optimal actions at each history.

To show that $T(P)$ is convex-valued, for any pool composition P consider two other pool compositions $P_1, P_2 \in T(P)$ generated by best response strategies $\sigma_1, \sigma_2 \in BR_+(P)$, and let $\tilde{P} \equiv \alpha P_1 + (1-\alpha) P_2$ be a convex combination of these two pools. Let $\tilde{\sigma} \equiv \alpha\sigma_1 + (1-\alpha)\sigma_2$ denote a strategy by which an agent initially randomizes between σ_1 and σ_2 with probabilities α and $1-\alpha$ and then follows the realized strategy thereafter. First, observe that by construction $\tilde{\sigma}$ is a best response strategy. Furthermore, $\tilde{\sigma}$ generates pool $\tilde{P}$ since $\Pr(a, \theta|\tilde{\sigma}, P) = \alpha \Pr(a, \theta|\sigma_1, P) + (1-\alpha)\Pr(a, \theta|\sigma_2, P) = \tilde{P}(a, \theta)$. Therefore, $\tilde{P} \in T(P)$.

To prove $T(P)$ is closed, we will argue that the extreme points of $T(P)$ are generated by pure strategies, and that in fact the set $T(P)$ is then the convex hull of these extreme points. To this end, define $S(P) \subset BR_+(P)$ as the set of pure strategy best responses. Given that $BR_+(P)$ is non-empty, so too is $S(P)$ since at any h at which a $\sigma(h)$ is an optimal mixed action, putting full weight on any of the pure actions must also be optimal.

For a fixed pool P let $\bar{Q}(a, \theta) \equiv \sup_{T(P)} Q(a, \theta)$ be the extreme point of the set $T(P)$ in the (a, θ) direction. Since $Q(a, \theta)$ is linear in $\sigma(a, h)$ at every history h for every pure action a, there exists a pure strategy $s \in S(P)$ so that $Q(a, \theta|s) = \bar{Q}(a, \theta)$. In other words, the extreme points of $T(P)$ are achieved using best response pure strategies. Then, since we previously showed $T(P)$ is convex, $T(P)$ is the convex hull of these extreme points and is therefore also closed. ■

Now, recall that $u(a, \mu(h|P)|P)$ is the expected stage payoff for an agent with a history h facing pool P and taking stage action a. Let $U(\sigma, \mu(h|P), \tau|P)$ be the expected continuation payoff (including the contemporaneous stage payoff) of an agent from following strategy $\sigma \in \Sigma_+$ when facing pool P and having history h with at most τ periods remaining, and let $U^*(\mu(h|P), \tau|P) = U(\sigma^*(P), \mu(h|P), \tau|P)$ be the agent's maximized payoff.

Claim 4 *The payoff to taking stage action a and optimizing thereafter* $V(a, \mu(h|P)|P) \equiv u(a, \mu(h|P)|P) + (1-\delta)E[U^*(\mu(h'|P), \tau - 1|P) \mid a, h]$ *is continuous in P.*

Proof of Claim: First, we show the belief $\mu(h|P)$ is continuous in P. This follows from the fact that for any particular stage observation $y_t = (s_t, a_t, q_t, \gamma_t)$, the likelihood of this observation given the agent is of type θ_i is given by equation (8), and is continuous in each $P(a, \theta)$ and thus in P. Therefore, the likelihood ratio of this given stage observation, and in turn the likelihood ratio of the sequence of stage observations $h = (y_1, ..., y_t)$, and thus the belief $\mu(h|P)$ are all continuous in P.

Then, every stage payoff $u(a, \mu(h|P)|P)$ is also continuous in P, since by equation (7) $u(\cdot)$ is continuous in its arguments μ and in P. This implies that the continuation value $U(\sigma, \mu(h|P), \tau|P)$ of following strategy σ is continuous in P, since it simply integrates stage payoffs.

In turn, the continuity in P of the continuation value $U(\sigma, \mu(h|P), \tau|P)$ for a given strategy σ implies the continuity in P of the maximized continuation value $U^*(\mu(h|P), \tau|P)$. That is, fix a history h and a pool P, let $\sigma^*(P)$ be an optimal strategy when the pool is P and

$\tilde{\sigma} \equiv \lim_{Q \to P} \sigma^*(Q)$ be the limit optimal strategy as the pool approaches P. If

$$\lim_{Q \to P} U(\sigma^*(Q), \mu(h|Q), \tau|Q) > U(\sigma^*(P), \mu(h|P), \tau|P)$$

then $U(\tilde{\sigma}(P), \mu(h|P), \tau|P) > U(\sigma^*(P), \mu(h|P), \tau|P)$ by the continuity of $U(\sigma, \mu(h|P), \tau|P)$ in P and μ. That is, the limit strategy $\tilde{\sigma}(P)$ does strictly better than the best strategy $\sigma^*(P)$ which is a contradiction. Conversely, if

$$\lim_{Q \to P} U(\sigma^*(Q), \mu(h|Q), \tau|Q) < U(\sigma^*(P), \mu(h|P), \tau|P)$$

then by the same continuity argument there will exist a Q sufficiently close to P so that $U(\sigma^*(P), \mu(h|Q), \tau|Q)) > U(\sigma^*(Q), \mu(h|Q), \tau|Q))$, again a contradiction. Therefore, $U^*(\mu(h|P), \tau|P)$ is continuous in P.

Finally, the probability of observing history h' conditional on today's history h and action a is derived from the likelihood function in (8) and is also continuous in P, thus making $E[U^*(\mu(h'|P), \tau - 1|P) \mid a, h]$ continuous in P. Therefore, since both the stage payoff $u(a, \mu(h|P)|P)$ and expected maximized continuation payoff are continuous in P, so is $V(a, \mu(h|P)|P)$. ■

Claim 5 *$T(P)$ is upperhemicontinuous.*

Proof of Claim: Toward a contradiction, suppose there exists a sequence $(P_i, Q_i) \to (P, Q)$ in which $Q_i \in T(P_i)$ for all i but $Q \notin T(P)$. We will argue that a consequence of this is an optimal strategy at P which gives a strictly higher payoff than best response strategies at nearby P_i, which is a contradiction.

Let $B(Q, \varepsilon)$ be an ε-neighborhood of pool Q and let $\Sigma_+(Q, \varepsilon|P)$ be the set of strategies that generate a pool distribution in $B(Q, \varepsilon)$ when facing pool P. Since $Q \notin T(P)$ and $T(P)$ is closed, there exists some $\bar{\varepsilon} > 0$ so that for $\varepsilon < \bar{\varepsilon}$ if $\sigma \in BR_+(P)$ and $\sigma' \in \Sigma_+(Q, \varepsilon|P)$ then there exists a $\delta(\varepsilon) > 0$ so that $V(\sigma|P) - V(\sigma'|P) > \delta(\varepsilon)$. In other words, when facing pool P any strategy that generates a pool close to Q is strictly suboptimal.

Now, observe that when facing a pool P_i sufficiently close to P, if σ_i generates Q_i when facing P_i then $\sigma_i \in \Sigma_+(Q, \varepsilon|P)$. That is the best response strategy σ_i when facing pool P_i would generate a pool in $B(Q, \varepsilon)$ when facing nearby pool P, which follows since $m(h|\sigma, P)$ is continuous in P as seen in equation (9).

Since $\sigma_i \in \Sigma_+(Q, \varepsilon|P)$ whenever P_i is sufficiently close to P, it follows that $V(\sigma|P) > V(\sigma_i|P)$. Then, by the continuity of $V(\sigma|P)$ in P established in Claim 4, it also follows that $V(\sigma|P_i) > V(\sigma_i|P_i)$ for P_i sufficiently close to P. This is a contradiction since it implies σ_i is not optimal at P_i. ∎

Putting together the results of the previous claims, since $\mathcal{P}$ is a non-empty, compact and convex subset of $\mathbb{R}^5$, and since the mapping $T : \mathcal{P} \to \mathcal{P}$ is upperhemicontinuous (Claim 5) non-empty, convex-valued, and closed (Claim 3), it follows by Kakutani's theorem that T has a fixed point. ∎