



ELSEVIER

Journal of Financial Markets 2 (1999) 153–178

Journal of
FINANCIAL
MARKETS

www.elsevier.nl/locate/econbase

Learning from others, reacting, and market quality¹

Rajesh Chakrabarti*, Richard Roll

The Anderson School at UCLA 110 Westwood Plaza, Box 951481 Los Angeles, CA 90095-1481, USA

Abstract

Traders pay attention to one another but are unable to perfectly deduce each others' beliefs from transactions alone. This explains why markets are hard to beat and also why trading occurs at all. Even when traders react rationally to the actions of others, they cannot arrive easily at a common posterior assessment of value. We model a realistic market composed of traders who combine their own private information with rational learning about the information possessed by others. We compare phenomena in this market with an otherwise identical market populated by traders who receive the same private information but ignore other traders. Using simulation to engender greater realism, we find that learning usually reduces volatility, increases the accuracy of the market price as a forecast of value, reduces trading volume, and decreases the prevalence of bubbles. However, for some combinations of market conditions, learning can have the opposite effect. The marginal influences of eight different market conditions, ranging from information heterogeneity through resource diversity, are estimated. Prices, volatility, volume, and bubbles exhibit subtle and complex responses to market conditions. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: G10; D83

Keywords: Learning; Incomplete information; Volatility; Market Quality

* Corresponding author. Tel.: 310 825 6118; fax: 310 206 8404; e-mail rchakrab@anderson.ucla.edu.

¹ Without blame, many thanks for comments and suggestions received on an earlier version of the paper from Antonio Bernardo, Blake LeBaron, Olivier Lédot, Bruce Lehmann, Tai Ma, David Mayers, Mark Stohs, an anonymous referee and participants in seminars at the Asia/Pacific Finance Conference, California State University Fullerton, the National Bureau of Economic Research, the University of California-Riverside, and National Sun-Yat Sen University.

1. Introduction

Visitors to exchanges and trading floors are often struck by the keen interest traders show in the actions of other traders. The decibel level is produced not just by orders and news, but also by loud discussions about the intentions of those thought to possess privileged information or market power. Scholars have long been intrigued by such behavior and have offered interpretations ranging across a wide spectrum of possibilities.

At one extreme, a famous passage in Keynes' (1937) *General Theory* expressed his view that traders care for nothing *except* each others' opinions. His 'musical chairs' portrait of speculative markets (pp. 155–157) disjoins price from long-term value and supposes volatility to be generated by trading itself, not by new information. In a Keynesian asset market, orders are determined exclusively by the observed or anticipated transactions of other participants, however misguided they may be. A trader buys or sells because someone admired has just done the same. Allegedly, such resolute mimicking induces high volatility and drifting prices unconstrained by rational anticipations of cash flows.

A diametrically opposite view is exemplified by the 'no trade theorems' of scholars such as Aumann (1976), Tirole (1982), or Milgrom and Stokey (1982). In their work as in Keynes', traders are fascinated by the opinions of others, but a far different implication is reached. Instead of reacting wildly, rational Bayesian traders assess the possible information possessed by others *before* placing their own orders. In the limit, given frictionless access to each others' probability beliefs, so-called 'common knowledge', trading ceases entirely because everyone comes to agree on the asset's value, even though initially they held divergent beliefs. As Hirshleifer (1977) puts it, if a 'fully complete market... were available, no one would need to use it', (p. 978), after an initial round of trading.

Casual empiricism reveals that neither extreme depicts reality. We know that most financial markets are more rational than Keynesian because profitable speculation is hard to achieve. We know too that financial markets exhibit a veritable blizzard of persistent trading activity, evidence of costs or other barriers along the path to common knowledge. The task facing current research is to erect models of speculative markets which are both realistic and consistent with the best efforts of rational agents.

Recent papers (e.g., De Long et al., 1990; Welch, 1992) have emphasized that *circumscribed* learning, deriving information by observing the actions of other traders without being able to fully infer their complete information sets, can lead to 'cascades' or 'speculative bubbles' under some circumstances, even if everyone is a good Bayesian. For example, Welch stipulates that traders act sequentially (barrier to complete knowledge #1.) Every trader observes whether previous traders have bought or sold, but no one is allowed to trade more than a single unit (barrier #2). Cascades then arise. The price might go to zero, even for an

asset with positive economic value, while a worthless asset might trade at a positive price. Yet everyone is rationally analyzing available information.

Welch's particular barriers to complete knowledge help an instructive example remain mathematically tractable. More realistic barriers would, of course, be desirable but they may necessitate other analytic apparatus such as numerical methods and simulation. This paper presents a simple but slightly more realistic model and presents conclusions numerically obtained.

In our model, as in Welch's, trading proceeds sequentially. Traders do not stand around and swap information about probability beliefs before acting. This 'barrier' is realistic. Trading in many financial markets proceeds more or less sequentially and organized exchanges with continuous outcry are exclusively sequential.² Since transactions occur one after the other, any agent is free to observe the entire sequence of previous prices and quantities. Whether *some* traders have *additional* information about previous traders is a key question and an important component of our research design.

Scharfstein and Stein (1990), Shiller (1990,1995) have suggested that traders keep track of one another. While Scharfstein and Stein lay emphasis on principal-agent relations to motivate such behavior, Shiller mentions an 'information-based' explanation. Supposedly, institutional investors and other high net worth players try to infer private information from the actions of competitors. Each would prefer maximum information about the asset being traded, but none is certain that his own information is reliable nor that it constitutes an adequate basis for trading. Everyone realizes that others eagerly seek information that, when found, is likely to precipitate action. This is why it makes sense to watch trading closely.

The identity of large block traders has momentous significance. Rumors circulate around sizable trades and other trades are often triggered, particularly if the initial trader is someone potentially knowledgeable, such as an insider. Similarly, fearing that other traders will run in front, the possessor of material private information often attempts to conceal his identity and to space orders temporally.

This suggests that a realistic model should permit heterogeneity across traders in the possession of information, the willingness to disclose whatever information is possessed, and the command over resources. Reality dictates also that such attributes be only partially inferable from market activities. Our simulation model possesses these features. We investigate the influence of reactive learning behavior on asset price accuracy and volatility. Tentative answers are suggested to the following questions: Do asset prices swing more when traders learn from each other or when they rely only on private

² Some continuous trading markets open with an initial Walrasian type auction; this may be better suited to the mutual discovery of beliefs.

information? What conditions increase or reduce the influence of reactive learning behavior on the amplitude of price fluctuations? How does trading behavior affect the quality of price as a forecast of economic value?

We specify a model of reactive learning, simulate trading, then compare the resulting prices to those observed when traders behave independently. With the same distributions of private information and wealth, we compute the excess volatility and trading volume, if any, induced by the extent of inter-reaction. We also assess the relative quality of prices, measuring how accurately prices forecast true economic value when traders react to each other and when they do not.

Section 2 of the paper describes our model. Section 3 presents the simulation results and offers interpretations. Section 4 concludes with suggestions for extensions in future research.

2. The model

2.1. General description

We model trade in a market for a single asset in fixed aggregate supply. The asset makes a liquidation cash payoff V , unknown during trading, which depends on the realizations of ‘fundamental’ variables. Each informed trader secures private information about one or more fundamental variables and therewith appraises the present value of V . For simplicity of illustration, trading occurs over a time interval short enough that discounting can be ignored.

The financial market is populated by two types of investors—numerous ‘small’ players and a few ‘big’ ones. A ‘big’ player is defined as someone *with enough resources* to buy or sell a quantity that attracts the attention of others, while a ‘small’ player’s trades go unnoticed. In our model, a particular trade is noticed if it involves at least a given number of units,³ a threshold we call a Minimum Observable Transaction ‘MOT’. So, by definition, a ‘big’ player can buy/sell at least a ‘MOT’ while a ‘small’ player cannot. A MOT-sized trade reveals something about the private signal of the big player; consequently, big players avidly follow and carefully analyze the observable trades of their fellows.

Small players could be regarded as liquidity traders or amateur speculators who, perhaps owing to their limited resources, do not find it worthwhile to keep track of the transactions of anyone, either big or small. In aggregate, small players provide a base excess supply function to big players. We assume an inverse supply function $p: \mathbf{R} \rightarrow (\text{LPP}, \text{HPP})$, positively sloped, and asymptotic to two bounds, the ‘highest plausible price’, HPP, and the ‘lowest plausible

³ A ‘unit’ would be, for example, a share of stock or a bond.

price', LPP, with the small players' market clearing at a price equal to the average of these bounds. We employ a particular function⁴ with these properties,

$$P = \frac{1}{2}(\text{HPP} + \text{LPP}) + [(\text{HPP} - \text{LPP})/\pi]\tan^{-1}q,$$

where P is the asset price and $q \in \mathbf{R}$ is the number of units bought ($q > 0$) or sold ($q < 0$) *cumulatively* by big traders. By trading a MOT or more, each big player adds to q , raising or lowering price to the value given by this inverse supply function.

Time is separated into several 'signal periods'. During each period, some big players receive private information about the future value of the asset, but no one becomes perfectly informed. Hence big players have limited faith in their own price forecasts. New information becomes available to all or to a subset of big players in each successive signal period.

Trading progresses sequentially, each big player moving after observing the moves of all previous players. After trading once, each big player joins the queue at the end and has another chance to move in the next trading 'round' within the same 'signal period'. Trading is transparent in that every big player observing a MOT-sized transaction also knows the perpetrator's identity.

Big players are endowed initially with differing amounts of cash *and* units of the asset. Neither short-selling nor borrowing is permitted. Hence, when a player buys or sells and stops at a certain price, others cannot be sure whether (a) he ran out of money or asset units, (b) the price reached his signal, or (c) he is attempting to mislead others by engaging in strategic behavior. Others are obliged to assess the relative likelihood of these possibilities after each observed trade.

We measure the result of their deliberations by a posterior distribution of value, V , derived by each trader by combining his private information with the information deduced from observing others. Traders are rational in the sense that Bayes' Rule is employed in deriving their respective posteriors. However, because the information of other traders is not perfectly observable, traders act without coming to complete agreement; their respective posteriors are diverse.

2.2. Formal specification

In each signal period, some big traders receive normally-distributed signals⁵ about value,

$$S_i \sim \mathcal{N}(\theta, \sigma^2),$$

⁴ This specification clearly injects an unrealistic element into our model. In its defense are simplicity and the fact that it would be easy to generalize.

⁵ The symbol $\sim$ denotes 'distributed as' and $\mathcal{N}(\cdot, \cdot)$ denotes the normal distribution.

where S_i is the signal received by trader i , Θ is the expected signal in a particular signal period and σ is the standard deviation across traders.⁶ Some big traders receive no signal.

Every trader possesses the same prior distribution on the mean signal, Θ , a ‘natural conjugate’ Normal

$$\Theta \sim \mathcal{N}(\mu, \tau^2)$$

whose parameters are common knowledge. For simplicity, we also assume that the volatility of signals, σ , is common knowledge and the same across traders, assumptions that can be easily relaxed in subsequent work.

The final value of the asset, V , is assumed to be a weighted average of the expected signal, Θ , and random noise, ζ ,

$$V = \kappa\Theta + (1 - \kappa)\zeta,$$

where the $0 < \kappa < 1$ is a constant and ζ is a completely unobservable Normal variate with mean Θ and variance *unknown* to traders. Clearly, $E(V) = \Theta$. Hence, by calculating posterior distributions on the expected signal Θ , traders are also deriving the best possible estimate of final value, recognizing that they are risk neutral and that the noise term ζ is unknowable.

The parameter κ , which we call ‘extractability’, measures the maximum possible information that all traders could derive about final value if they were to share information completely and agree on a common posterior. This contrivance is adopted in an effort to measure the impact of reacting and learning under varying degrees of maximum *possible* information aggregation. It is conceivable a priori, for example, that cascades and bubbles arise with greater frequency when perfectly aggregated information is at best relatively poor (or vice versa).

To understand the Bayesian updating rule, first consider a trader who ignores the actions of others. This is necessarily the status also of the very first trader. Combining his private signal with the common prior, his posterior on expected value is

$$\pi(\Theta|S_1) \sim \mathcal{N}[\mu(S_1), \rho_1^2],$$

where

$$\rho_1^2 \equiv \rho_{\text{NR}}^2 = \tau^2\sigma^2/(\tau^2 + \sigma^2),$$

⁶ Since the Normal is unbounded, this formulation theoretically permits implausible signals, such as values below zero. We finesse this difficulty by choosing parameters which produce such implausible results only rarely. Note, however, that negative, or very large *values* are theoretically possible, though *prices*, (as opposed to *values*), are bounded above zero and below a maximum plausible level.

and

$$\mu(S_1) = (\mu/\tau^2 + S_1/\sigma^2)\rho_1^2 = [\sigma^2/(\tau^2 + \sigma^2)]\mu + [\tau^2/(\tau^2 + \sigma^2)]S_1.$$

See Berger (1985), pp. 127–128. This is also the form of the posterior for all traders who pay no attention to others; the only difference would be that each trader j 's private signal S_j replaces S_1 , the first trader's signal, in the formula. Since by assumption τ and σ are common knowledge, the volatility ρ_1 will be identical in posteriors of traders who have No Reaction (NR) to others, hence we denote it ρ_{NR} . It appears also as one element of posterior volatility when traders do React to Others (RO).

If everyone's signal were observable without error, the sample mean of all signals is a sufficient statistic for the population mean. In Bayesian analysis, a sufficient statistic can be substituted in Bayes' Rule for *all* of the individual sample observations,⁷ Consequently, with $\bar{S} \equiv \sum_j S_j/N$ for N traders, the *common* posterior would be

$$\pi(\Theta|\bar{S}) \sim \mathcal{N}[\mu(\bar{S}), \rho_N^2],$$

where $\rho_N^2 = \tau^2\sigma^2/(N\tau^2 + \sigma^2)$ and $\mu(\bar{S}) = [(\sigma^2/N)/(\tau^2 + \sigma^2/N)]\mu + [\tau^2/(\tau^2 + \sigma^2/N)]\bar{S}$. The signals in our model are not observed without error, so no trader derives this common posterior. We record it here for later comparison with the actual posteriors, which diverge from each other.

2.2.1. Incompletely observable signals

Let P_{j-1} denote the prevailing market price, i.e., the price just before trader j decides to act. Let θ_j denote the mean of trader j 's posterior distribution of Θ based on his own signal and the actions of all previous traders. Assume also that every trader, by acting, provides noisy information about his own private signal according to the following scheme: After he finishes trading, the new price, P_j , satisfies

$$P_j = \theta_j - \xi,$$

where the random noise follows

$$\xi \sim \mathcal{N}(\alpha_B, \gamma^2) \text{ for buyers}$$

$$\xi \sim \mathcal{N}(\alpha_S, \gamma^2) \text{ for sellers}$$

with γ known and $\alpha_S < 0 < \alpha_B$.

⁷ Cf. Berger (1985), Lemma 1, p. 127. An early application similar to our set-up is in Grossman (1976) who shows that price itself is a sufficient statistic for information in some Bayesian updating circumstances.

Although the trading noise term ξ is modeled as a random variable, we are really imagining it as a choice made by each individual trader, a choice unknowable and hence random from the perspective of other traders. Several distinct trading proclivities are covered by this single noise specification. First, a trader could be acting strategically; in an effort to mislead competitors, he might cease buying (selling) before the price rises (falls) to his current posterior estimate of value. Second, recognizing the winner's curse, a trader knows that his own posterior estimate is not perfect and may wish to bias his buying or selling toward some perceived plausible value. Third, each trader has limited resources and might stop buying simply because those resources are depleted.

These arguments imply that the reservation price of a buyer (seller) should be biased downward (upward) relative to the mean of his posterior distribution of value; hence, the biases $\alpha_B > 0$ and $\alpha_S < 0$. Bias alone does not, however, capture the full essence of incomplete signal observability. There must also be some heterogeneity in behavior across traders induced by diversity in resource endowments, strategic trading, etc. We parameterize these concepts with randomness in the trading noise.

The second trader is the first to witness another trader's action. He deduces that the distribution of the first trader's signal is

$$S_1 \sim \mathcal{N}(\Theta + \alpha, \sigma^2 + \gamma^2).$$

Consequently, the second trader's posterior distribution of value *taking account of the first trader's action* only is

$$\pi(\Theta|P_1) \sim \mathcal{N}[\mu(P_1), \rho_{RO}^2],$$

where

$$\rho_{RO}^2 \equiv [\tau^2(\sigma^2 + \gamma^2)]/(\tau^2 + \sigma^2 + \gamma^2) \text{ and}$$

$$\mu(P_1) = [(\sigma^2 + \gamma^2)/(\tau^2 + \sigma^2 + \gamma^2)]\mu + [\tau^2/(\tau^2 + \sigma^2 + \gamma^2)](P_1 + \alpha).$$

Note that α is either α_B or α_S depending on whether trader #1 buys or sells.

Combining his posterior distribution of the first trader's signal with his own private signal using Bayes Rule, the second trader calculates his *own* posterior distribution for value as

$$\pi(\Theta|P_1 \cap S_2) \sim \mathcal{N}[\theta_2, \rho_2^2],$$

where $\rho_2^2 = \rho_{RO}^2 \rho_{NR}^2 / (\rho_{RO}^2 + \rho_{NR}^2)$ and $\theta_2 = [\rho_{NR}^2 / (\rho_{RO}^2 + \rho_{NR}^2)](P_1 + \alpha) + [\rho_{RO}^2 / (\rho_{RO}^2 + \rho_{NR}^2)]S_2$. Trader #2 is motivated to buy if this posterior mean is above P_1 and to sell otherwise. His decision rule is: stop trading when the new price P_2 reaches $\theta_2 - \xi$, where ξ is a single draw from the trading noise distribution. Note that he will not buy (sell) at all if $\theta_2 - \xi$ turns out to be less (greater) than the prevailing price P_1 .

It is possible, but tedious, to solve the third and subsequent traders' problems recursively, each trader combining the actions of all previous traders with his own signal. However, a moment's reflection reveals that the previous trading prices contain all the information that any trader has about the signals received by others. The mean trading price, adjusted for the bias in the signal noise distribution, is therefore a sufficient statistic for all previous players' information. Exploiting the previously mentioned result from Bayesian analysis that a sufficient statistic can substitute for all of the individual observations, the posterior of each trader can be expressed solely in terms of the trader's private signal and the mean bias-adjusted price over previous trades.

Suppose, then, that a later trader, j , has witnessed N_B buyers and N_S sellers that have traded at prices P_{Bi} , $\{i = 1, \dots, N_B\}$ and P_{Sk} , $\{k = 1, \dots, N_S\}$. (The total number of active preceding traders is $N = N_B + N_S$.) Denote the means of buying and selling prices, respectively, as

$$\bar{P}_B = \sum_i P_{Bi}/N_B \text{ and } \bar{P}_S = \sum_k P_{Sk}/N_S.$$

A sufficient statistic is the bias-adjusted mean,

$$\mathbf{P} \equiv [N_B(\bar{P}_B + \alpha_B) + N_S(\bar{P}_S + \alpha_S)]/N.$$

Consequently, the posterior distribution of value derived by trader j is

$$\pi(\Theta|\mathbf{P} \cap S_j) \sim \mathcal{N}[\theta_j, \rho_j^2],$$

where

$$\rho_j^2 = \rho_{NR}^2 \rho_P^2 / (\rho_{NR}^2 + \rho_P^2),$$

$$\theta_j = [\rho_{NR}^2 / (\rho_{NR}^2 + \rho_P^2)] \mu_{(\mathbf{P})} + [\rho_P^2 / (\rho_{NR}^2 + \rho_P^2)] S_j,$$

$$\rho_P^2 = [\tau^2(\sigma^2 + \gamma^2)/N] / [\tau^2 + (\sigma^2 + \gamma^2)/N],$$

and

$$\mu(\mathbf{P}) = \{[(\sigma^2 + \gamma^2)/N] / [\tau^2 + (\sigma^2 + \gamma^2)/N]\} \mu + \{\tau^2 / [\tau^2 + (\sigma^2 + \gamma^2)/N]\} \mathbf{P}.$$

As trading progresses, and N increases, the posterior becomes ever more heavily influenced by the mean bias-adjusted price $\mathbf{P}$. Less reliance is placed on both the trader's own signal and on the prior mean before the first trade. Nonetheless, trading will not cease after the first round and will not die out completely until N becomes so large that all traders possess essentially the same posterior about Θ .

Terms and definitions are presented in the glossary of Table 1 for easy reference.

Table 1
Glossary of terms

MOT	Minimum observable transaction (in asset units)
Big player	One who can carry out a MOT
V	Economic value (after resolution of uncertainty)
HPP	Highest plausible price
LPP	Lowest plausible price
Signal	Private information about value
Θ	Expected value of V = Expected signal
Extractability, κ	Aggregate quality of signals; $V = \kappa\Theta + (1 - \kappa)$ (noise)
Clueless	Proportion of big players who receive <i>No</i> signal
Signal diversity, σ	Cross-trader volatility of signals
Trading noise, γ	Volatility rendering signals incompletely observable
$\alpha_B > 0$	Trading bias of buyers; i.e., expected value less reservation price
$\alpha_s < 0$	Trading bias of sellers

3. Simulations

3.1. Methods

Our simulations⁸ have two basic scenarios – (RO) when big traders React to Observable trades, combining private information with signals inferred from the actions of others according to the method previously described, and (NR) when there is No Reaction to the trades of others, everyone basing decisions solely on privately-received information.

Within each scenario, sequences of simulated prices and returns are used to compute a percentage ‘excess volatility’ statistic,

$$\omega_i \equiv [200(\sigma_{i,RO} - \sigma_{i,NR})/(\sigma_{i,RO} + \sigma_{i,NR})], \quad (1)$$

where $\sigma_{i,S}$ is the standard deviation across all simulated trades within replication i and scenario S .⁹ These scenarios differ only in that private beliefs are revised by learning under RO. Wealth and private signals are the same in either case. Clearly, $\omega_i > 0$ implies that reacting to others induces higher volatility while $\omega_i < 0$ means that it reduces volatility. The average (across T replications)

⁸ The code is available on request.

⁹ At first glance, it may seem unusual to make a comparison between x and y by calculating their difference $(x - y)$ divided by their average $\frac{1}{2}(x + y)$. But since both x and y are subject to sampling error, using either one as a base (denominator) in the comparison is subject to the difficulty that it could be vanishingly small just by happenstance. Hence, a random outlier could acquire a spurious importance. Our statistic reduces the chance that this will occur. The normalizing constant $200 = 100/(\frac{1}{2})$ simply assures that the statistic lies between $+100\%$ and -100% .

excess volatility statistic, $\varpi \equiv \sum_i \omega_i / T$ will be reported below for prices and returns, distinguished with mnemonic subscripts; i.e., ϖ_p and ϖ_r , respectively.

The excess volatility of returns requires no motivation, but some may be surprised that we also report the excess volatility of price. Our reasoning is that cascades or bubbles might take the form of *steady* price drifts with relatively low return volatility. For example, imagine a bubble consisting of a constant positive return at each transaction. The price volatility would be large (around its sample mean) while the return volatility would be zero.

We also investigate whether learning from and reacting to others helps the market price approach value. An ‘average valuation mistake’, is defined as:

$$M \equiv 200 \{ [\sum_i (\varphi_{i, \text{RO}} - V_i)^2]^{1/2} - [\sum_i (\varphi_{i, \text{NR}} - V_i)^2]^{1/2} \} / \{ [\sum_i (\varphi_{i, \text{RO}} - V_i)^2]^{1/2} + [\sum_i (\varphi_{i, \text{NR}} - V_i)^2]^{1/2} \}, \quad (2)$$

where $\varphi_{i, S}$ is the *final* price at the end of trading under scenario S and replication i of the simulation, and V_i is the economic value revealed after trading in replication i . Essentially, M is a relative¹⁰ root mean square prediction error over replications.

Learning could also influence the trading volume. We investigate this issue by quantifying the excess dollar volume, $\mathcal{V}$, analogously to ω_p and ω_r ,

$$\mathcal{V}_i \equiv [200(v_{i, \text{RO}} - v_{i, \text{NR}}) / (v_{i, \text{RO}} + v_{i, \text{NR}})], \quad (3)$$

where $v_{i, S}$ is the dollar volume of trade in replication i and scenario S .¹¹ The mean excess dollar volume over T replications is just $\sum_i \mathcal{V}_i / T$.

Speculative bubbles are most often defined in the literature as large excursions by the market price away from fundamental value. In simulated markets, there is a simple and readily measurable counterpart to the intuitive notion of a bubble; it is the difference between price and the true expected value of the asset. This difference can be calculated for every transaction price, so it seems reasonable to measure a ‘pseudo-bubble’ by the maximum absolute difference during trading,

$$\beta_{i, S} \equiv \max_t |P_{i, S, t} - \Theta_i|,$$

where t covers transactions observed during replication i and scenario S ; $\Theta_i = E(V_i)$ varies only across replications.¹² The impact of learning on the

¹⁰ Relative across the two basic scenarios.

¹¹ We have also measured volume in units of shares rather than dollars. The results are similar.

¹² The intuitive notion of a bubble relates to a large deviation between price and value, but since ‘large’ has never been specified by any authority, we simply record the maximum excursion, whatever it may be.

propensity for markets to depart on bubble trajectories can then be measured by an excess bubble statistic, defined as

$$\mathcal{B}_i \equiv [200(\beta_{i,RO} - \beta_{i,NR})/(\beta_{i,RO} + \beta_{i,NR})]. \tag{4}$$

In general, each market phenomenon may or may not depend upon market conditions. This is what we are seeking to ascertain. Market conditions we vary across simulations include the minimum noticeable trade size, or ‘MOT’; the diversity of the signals across players; the aggregate quality of all signals, ‘extractability’; and the extent of trading noise induced by strategic trading or endowments, (of both cash and assets). We also investigate the interactions of market parameters and document whether phenomena such as bubbles are more likely to occur under certain combinations.

3.2. *Parameters, (i.e., market characteristics)*

Average excess price and return volatility and average valuation mistake were computed for all combinations of the parameters values listed in Table 2.

The price at the beginning of every Signal Period (replication) is initialized to 100. Then trading begins by rounds, individual Big Traders following in the same order within each round. Since there are five trading rounds and twenty Big Players, there are 100 transactions in each simulated Signal Period. Price

Table 2
Market Conditions (Parameters) for simulated trading

<i>Parameters held constant</i>			
Highest plausible price (HPP)			200
Lowest plausible price (LPP)			0
Mean of the common prior on value, μ			100
Number of replications (signal periods)			30
Number of big players			20
Number of trading rounds per signal period			5
<i>Parameters varied</i>			
Volatility of prior, ^a τ	2.5	5.0	7.5
Signal diversity, σ	2.5	5.0	7.5
Trading noise, γ	1	2	3
Buyer bias relative to signal, α_B	1	2	3
Seller bias relative to signal, α_S	− 1	− 2	− 3
MOT	0.0025	0.0075	0.0125
Clueless, %	0	25	50
Extractability	0.25	0.5	0.75

^aJust to keep things manageable, in all simulations we set the volatility of aggregate information noise, ζ equal to the volatility of the prior.

and return volatilities are computed from these 100 observations. The ‘valuation mistake’ or pricing error is computed at the end of each signal period by comparing the final market price to value; (Cf. The definition of ‘extractability’ in the Table 1 glossary.)

3.3. Simulation results: The influence of reactive learning on volatility

The excess price volatility (see Eq. (1)), $\bar{\sigma}_p$, averaged across all parameter values, was -28% , i.e., reactive learning lowered price volatility substantially on average over our parameter space. The overall average of the excess return volatility $\bar{\sigma}_r$, was similar, -30% . Within the artificial setting of these simulations, trader interaction appears to stabilize market prices and returns, at least on average for the chosen parameter values.

There is, however, considerable variation in this stability gain across different market conditions (parameter sets). This is partly attributable to randomness, of course, but also partly due to the market conditions themselves. The maximum observed $\bar{\sigma}_p$ for any parameter set was 123% and the minimum was -121% . The maximum $\bar{\sigma}_r$ was 162% and the minimum was -139% . There is evidently some cross-replication asymmetry in the distribution of the excess return volatility statistic.

In accordance with ‘no trade’ literature, we noticed that the decreased return volatility was often associated with a reduction in the *number* of trades over time accompanied by a stabilized price and virtually zero returns across later transactions.

To summarize compactly the various influences of individual parameters, we report regressions with parameter values as explanatory variables. Ordinary least squares (OLS) regressions are reported in Table 3, whose panels are for different dependent variables: average excess price volatility $\bar{\sigma}_p$ in the upper panel, and average excess return volatility $\bar{\sigma}_r$ in the lower. The explanatory variables, the parameters, vary as shown in Table 2 and the sample size is the total number of different unique parameter sets.¹³

It may seem a bit unusual to report simulation results by regressions of the simulation output on input parameters. In this case, however, the procedure seems warranted. Our simulated market is complex and is characterized by eight varying conditions along with some fixed conditions (Table 2). We have no a priori intuition about how these conditions should influence market attributes such as volatility and pricing error nor whether such influences have a particular functional form nor whether they have strong interactions. It would be exceedingly tedious to vary just one parameter at a time and observe the impact on market phenomena. Moreover, to the extent that parameters interact, this could

¹³ There are $3^8 = 6561$ different combinations of parameters.

Table 3

Excess volatility from reactive learning: marginal influences of market characteristics

Market attribute	Coefficient	T-statistic
<i>Dependent variable: average excess price volatility, $\bar{\sigma}_p$</i>		
Volatility of prior, τ	– 16.3	– 113.4
Signal diversity, σ	3.46	24.2
Trading noise, γ	3.39	9.45
Buyer bias, α_B	6.25	17.4
Seller bias, α_s	– 6.30	– 17.6
MOT	207	2.88
Clueless, %	0.108	7.54
Extractability	0.161	0.112
Adjusted R^2 : 0.684		
<i>Dependent variable: average excess return volatility, $\bar{\sigma}_r$</i>		
Volatility of prior, τ	– 20.2	– 109
Signal diversity, σ	6.89	37.3
Trading Noise, γ	13.8	29.9
Buyer bias, α_B	0.922	2.00
Seller bias, α_s	– 1.44	– 3.12
MOT	315	3.40
Clueless, %	0.501	27.1
Extractability	0.439	0.238
Adjusted R^2 : 0.695		

Note: In a linear regression, the dependent variable was the value of the average excess volatility statistic Eq. (1), observed for a given set of market conditions, (parameter values). The independent variables are the parameter values. The number of observations in the regression, 6561, is the number of different combinations of parameter values.

be quite misleading. Regressions offer an informative and concise device for reporting the essential features and the complexities of our simulated markets.

The observations in these regressions are independent by construction because different random numbers were drawn for each replication. However, the regressions might be mis-specified because of non-linearity or heteroskedasticity, two potential problems that we investigate below.

For each regression, we report *t*-statistics and adjusted *R*-squares, but one should be forewarned that these are somewhat arbitrary; presumably, they could be raised (perhaps to perfection with the correct non-linear functional form) simply by increasing the grid density of parameter values. The grid of parameter values listed in Table 2 is sufficiently dense to produce significant and meaningful inferences and yet sparse enough to be computationally tractable. We make no assertion that it is optimal.

The results in Table 3 for both price and return reveal that the reduction of volatility from learning is strongly enhanced by a more diffuse prior. The marginal influence of the prior's volatility, τ , is sharply negative. Learning is more effective in reducing market volatility when there is more to learn, i.e., when traders do not possess strong prior convictions and hence place greater emphasis on the actions of others. A diffuse prior implies also that they would weight their private signal more in the no learning (NR) case, but this is overwhelmed in the learning (RO) case by assigning greater credence to the actions of others.

Signal Diversity, σ , significantly increases volatility under learning (RO) relative to no learning, (NR). Similarly, a large component of trading noise, γ , seems to mislead good Bayesians attempting to learn from others; it too raises both price and return volatility under RO relative to NR. The same marginal effect can be observed for the buyers' and sellers' mean bias, (i.e., the difference on average between their posterior mean signal and their reservation price.) When these biases become larger in absolute value, other traders are misled enough to induce an increase in market volatility.

A smaller minimum observable transaction, or 'MOT', enhances the volatility reduction induced by learning. Even though traders are misled by others emitting false or volatile signals, they are astute enough to learn on average. Consequently, observing more individual events, (a smaller MOT), improves their ability to come to an agreement and thereby results in smaller price and return volatility. We observe a similar phenomena in the variable 'Clueless' which measures the fraction of traders who receive no signals at all. Increasing that fraction significantly worsens the volatility improvement under learning. So, even though some actions are misleading, it is better on average to observe more of them.

Finally, 'extractability', the average forecast quality of private signals, has no influence whatsoever on excess volatility. At first, this may seem puzzling. But prior to the final revelation of value, traders have no information about value other than their own private signals and what can be deduced from the actions of others. Their trading can depend only on this information. Reacting to others reduces volatility on average, but that reduction *cannot* be related to something unknowable, the aggregate signal quality.

3.4. *The quality of price as a forecast of value*

As in the case of excess price volatility and excess return volatility, a negative value of M , the average valuation mistake (Eq. (2)), indicates that players are (collectively) forecasting value better when they react to each others' trades than when they ignore the information embedded in those trades.

Under many different market conditions, reactive learning increases the forecasting quality of the final transaction price. Across all parameter values, the

overall average valuation mistake, M , was -12.2% , a material reduction in the root mean square forecasting error provided by reactive learning. The improvement was not confined to a few outliers; In over 78% of parameter combinations, the average valuation mistake was negative.

The marginal influences of the parameters on M is summarized by the OLS regressions in Table 4. In an interesting contrast with the volatility results reported in Table 3, the influences of a diffuse prior and of signal diversity are reversed. A less diffuse prior and larger cross-trader dispersion in private signals significantly *increase* the accuracy of the market price when learning (RO) takes place, relative to the pricing accuracy when there is no learning (NR).

This intriguing result was quite surprising, at least to us. Evidently, a diffuse prior diminishes volatility under learning or raises it under no learning (or both) while it also decreases the accuracy gained through learning. Similarly, greater signal diversity tends to exacerbate volatility under learning relative to no learning, but it enhances the ability of learning to increase accuracy.

These findings suggest that markets should not be judged by volatility alone. Indeed, a high level of volatility induced by traders reacting to each other may actually be associated with a more informative price. We see here, for instance, that greater heterogeneity in private signals, though causing more volatility, brings the market price at the close of trading closer to fundamental value. Similarly, the extent of the diffuseness of the common prior has opposite effects on volatility and pricing error.

Another contrast between Tables 3 and 4 is the influence of ‘extractability’. It had no influence on volatility (Table 3) but significantly reduces the pricing mistake (Table 4). This is intuitively plausible in that the average quality of

Table 4

The quality of price as a forecast of value: marginal influences of market characteristics

Market attribute	Coefficient	<i>T</i> -statistic
Volatility of prior, τ	0.364	4.30
Signal diversity, σ	-1.81	-21.4
Trading noise, γ	0.356	1.68
Buyer bias, α_B	1.46	6.91
Seller bias, α_S	-0.697	-3.30
MOT	91.8	2.17
Clueless, %	0.243	28.7
Extractability	-17.9	-21.2

Adjusted R^2 : 0.217

Note: In a linear regression, the dependent variable was the value of the average valuation mistake M , Eq. (2), observed for a given set of market conditions, (parameter values). The independent variables are the parameter values. The number of observations in the regression, 6561, is the number of different combinations of parameter values.

private signals is closer to the true value when extractability is large. Hence, reactive learning improves the market's ability to aggregate those higher quality private signals. Similarly, as one would have anticipated, the pricing mistake is lower when more agents receive signals and also react to each other.

The other market conditions such as endowment differences, 'Clueless', and 'MOT' have the same marginal effects that they had on volatility, though the significance levels are considerably smaller.

3.5. Trading volume

Across all parameter combinations, the excess dollar volume (Eq. (3)) averaged -37.5% , a substantial reduction in volume induced by reactive learning. However, there was significant variation across replications with a maximum 47.3% and minimum of -72.3% . The average, maximum and minimum for the excess share volume were approximately the same as the excess dollar volume. These results are compatible with the basic thrust of 'no-trade' theorems in the sense that reactive learning tends to diminish trading.

Marginal influences of market conditions, (i.e., of parameters) on excess dollar volume are provided by the regression reported in Table 5.¹⁴

Excess volume mirrors the behavior of volatility, which is strikingly apparent in comparing Tables 3 and 5. An empirical association between volatility and volume has long been documented in the literature; cf. Karpoff (1987) or Gallant, Rossi, and Tauchen (1992). It has generally been supposed that this connection arises when new information arrives and evokes trading; however, we see from these simulation results that an information event is not necessary. The volume/volatility connection can arise spontaneously from trader interaction.

There is nonetheless one parameter which seems to influence volatility differently than volume; viz., the trading biases (α 's) of sellers and buyers. As Table 3 reports, a bigger absolute bias decreases the volatility reduction induced by learning behavior. The signs of both α 's are reversed in the trading volume regression reported in Table 5, thereby indicating that greater bias increases the volume reduction from learning. Admittedly, the coefficients in Table 5 are only marginally significant.

We can think intuitively of these trading biases, the average difference between a trader's signal and his reservation price, as measures of endowments, *inter alia*. A larger bias implies that a typical buyer (seller) runs out of money (shares) farther away from his posterior assessment of the asset's value, before he has bought (sold) as many shares as his posterior would have dictated in the

¹⁴ Results for share volume are very similar and are not reported to save space. They will be provided to interested readers upon request.

Table 5
Marginal influences of market characteristics on volume

Market attribute	Coefficient	<i>T</i> -statistic
Volatility of prior, τ	– 26.5	– 113
Signal diversity, σ	9.56	40.9
Trading noise, γ	25.4	43.5
Buyer bias, α_B	– 1.03	– 1.76
Seller bias, α_s	1.28	2.18
MOT	470	4.02
Clueless, %	0.924	39.5
Extractability	0.334	0.143
Adjusted R^2 : 0.732		

Note: In a linear regression, the dependent variable was the value of the average valuation mistake $\mathcal{V}$, Eq. (3), observed for a given set of market conditions, (parameter values). The independent variables are the parameter values. The number of observations in the regression, 6561, is the number of different combinations of parameter values.

absence of a resource constraint. Evidently, the possibility of resource exhaustion interacts with learning to increase volatility but decrease volume. Again intuitively, the volume reduction arises because trading ceases sooner but volatility at the same time increases because the information conveyed to other traders is less precise.

3.6. The prevalence of bubbles

On average over all parameter values in our simulations, the excess prevalence of bubbles (Eq. (4)) had a mean value of – 13.0%, indicating a moderate reduction in bubbles induced by learning. There was, however, considerable variability across market conditions (parameters.) The standard deviation was 24% and the maximum and minimum were 82.0% and – 71.6%, respectively. Evidently, under some conditions, learning exacerbates the tendency for markets to form bubbles.

To shed some light on what particular market conditions might cause learning to exacerbate bubble formation, Table 6 records the marginal influences of various parameters. One might have anticipated that the results would be somewhat like the excess price volatility reported in Table 3 because greater excursions of price from value are likely also to represent greater departures of individual transaction prices from their mean. Indeed, this has occurred. Tables 6 and 3 have similar sign patterns and significance levels. The only sign difference is for Trading Noise, γ , which is negative in Table 6. Apparently, though trading noise and learning interact to produce more volatility, very large excursions from fundamental value, i.e., bubbles, are less likely when traders

Table 6
Marginal influences of market characteristics on bubbles

Market attribute	Coefficient	<i>T</i> -statistic
Volatility of prior, τ	– 8.21	– 89.6
Signal diversity, σ	2.95	32.2
Trading noise, γ	– 0.509	– 2.22
Buyer bias, α_B	2.43	10.6
Seller bias, α_s	– 2.30	– 10.0
MOT	142	3.09
Clueless, %	0.299	32.7
Extractability	0.394	0.430
Adjusted R^2 : 0.612		

Note: In a linear regression, the dependent variable was the value of the excess bubble statistic, $\mathcal{B}$, Eq. (4), observed for a given set of market conditions, (parameter values). The independent variables are the parameter values. The number of observations in the regression, 6561, is the number of different combinations of parameter values.

realize that they should be cautious in reacting when other traders are generating lots of noise.

3.7. Specification diagnostics

The regressions reported in the previous section are compact devices for presenting simulation results. They are not necessarily well-specified statistical models. A priori, there are at least two reasons to anticipate mis-specification. First, although the *observations* are independent by construction, possible non-linearity might have produced dependence in the *disturbances*. OLS fits a hyper-plane through the *observations*. If the true functional form has curvature, there might be dependence across successive disturbances when the observations are ordered by any one of the explanatory variables. Second, there is every reason to anticipate heteroskedastic disturbances. Both of these problems can be corrected, heteroskedasticity by econometric methods and non-linearity by estimating an appropriate functional form.

To correct for heteroskedasticity, we re-computed the previous regressions with the White (1980) heteroskedasticity-consistent disturbance covariance estimate. The results were virtually identical to the simple OLS results.¹⁵

Since we have no intuition about the extent and complexity of the hyper-surface curvature, if any, we utilize an agnostic investigative tool, a series

¹⁵ A copy of the tables will be furnished on request to interested readers.

expansion that can approximate any differentiable function. Let y denote the dependent variable and x_j the j^{th} explanatory variable related by an unknown surface $y = F(x_1, x_2, \dots)$. A Taylor series expansion is

$$y = F(\mu_1, \mu_2, \dots) + \sum_j (x_j - \mu_j) F_j + \frac{1}{2} \sum_j \sum_k (x_j - \mu_j)(x_k - \mu_k) F_{j,k} + \text{higher-order terms},$$

where $F_j \equiv \partial F / \partial x_j$ and $F_{j,k} \equiv \partial^2 F / \partial x_j \partial x_k$. To fit such an expansion empirically, terms higher than second order can be discarded and μ_j specified as the sample mean of explanatory variable j . In addition, each variable can be scaled by its sample standard deviation σ_j to help reduce multicollinearity, which might otherwise be a quandary.¹⁶ Scaled variates in regressions would then all have sample mean zero and sample standard deviation unity.

The regression analog of the Taylor series expansion is

$$(y - \mu_y) / \sigma_y = a + \sum_j b_j z_j + \sum_j c_j z_j^2 + \sum_j \sum_{k > j} d_{j,k} z_j z_k,$$

where a, b, c , and d are coefficients to be estimated and $z_j \equiv (x_j - \mu_j) / \sigma_j$. Estimated regression coefficients can be interpreted as functions of hyper-surface parameters; i.e.,

Intercept:

$$a = [F(\mu_1, \mu_2, \dots) - \mu_y] / \sigma_y.$$

Linear Terms:

$$b_j = (\sigma_j / \sigma_y) \partial F / \partial x_j.$$

Quadratic Terms:

$$c_j = \frac{1}{2} (\sigma_j^2 / \sigma_y) \partial^2 F / \partial x_j^2.$$

Cross-product Terms:

$$d_{j,k} = (\sigma_j \sigma_k / \sigma_y) \partial^2 F / \partial x_j \partial x_k.$$

The t -statistics¹⁷ from volatility surface estimation are reported in Table 7. Reassuringly, the linear terms have exactly the same pattern and higher levels of significance than their corresponding coefficients in the Table 3 regressions. This implies that surface curvature does not cause egregious mis-estimation of marginal linear influences, even when a simple plane is fit to the data. The

¹⁶ For instance, $(x_j - \mu_j) / \sigma_j$ and $[(x_j - \mu_j) / \sigma_j]^2$, though functionally dependent, are likely to be less co-linear than the raw variates x_j and x_j^2 .

¹⁷ To save space, we report only the t -statistics and adjusted R^2 s. All other results will be provided to interested readers. The regressions employ the White (1980) heteroskedastic consistent covariance matrix.

adjusted R^2 's have increased materially, an improvement to be expected when the surface is non-linear. Several parameters have highly significant quadratic terms and there are numerous strong interactions.

For both price and return excess volatility, Signal Diversity, σ , Diffuseness of the Prior, τ and Trading Noise volatility, γ , all have significant positive quadratic terms, thereby implying a concave downward surface near their respective means. These parameters also display significant interactions, negative between σ and both τ and γ , positive between τ and γ . A negative interaction implies a reduction in the slope of the surface when both parameters increase, and vice versa.

To give a concrete example, consider the influence of Signal Diversity on the excess volatility from learning. Its linear effect is positive, thus indicating that greater cross-trader heterogeneity in signals reduces the mitigating effect of learning on volatility. However, this linear effect is reduced when the prior distribution is more diffuse; i.e., the more uncertain traders are prior to seeing the actions of others, the smaller the marginal influence of Signal Diversity on the learning-induced reduction in volatility.

Another interesting example involves the interactions between buyer and seller biases, α_B and α_S , respectively, and trading noise volatility, γ . The two biases reduce learning-induced volatility reduction; their estimated linear effects show an increase in excess volatility with an increase in their absolute values. These marginal effects are both reduced by higher trading noise volatility. The cross-product term for α_B and γ (α_S and γ) is negative (positive). Understandably, less precise information about the trading biases mitigates their influence on learning.

Similar interpretations could be drawn out for other combinations of parameters. The general impression is a rather complex surface contouring the subtle influences and delicate interactions of market conditions. Although these results are intuitively plausible once they are disclosed by the simulations, it would probably have been difficult to conceive of them a priori.

Turning now to the non-linear aspects of the Average Valuation Mistake, M , the results in Table 8 reveal a slight improvement in explanatory power with the addition of non-linear explanatory variables. The adjusted R^2 has increased from 0.217 to 0.283.

Again, there are no surprises; the linear effects all have the same sign and slightly greater significance than the planar regression of Table 4. As with excess volatility, the Average Valuation Mistake surface is convex downward near the means of the volatility parameters, τ , σ , and γ . The interaction effects, however, are distinctly different.

In contrast to the volatilities, the important extractability parameter displays no significant non-linear (quadratic) effect. Increasing extractability augments the benefits from learning in forecasting value, (its linear coefficient is negative), but this effect is reduced, as might have been expected, by a more diffuse prior and by greater heterogeneity in private signals.

For excess volume and bubble prevalence, surface estimation again reveals non-linearity and frequent interaction effects.¹⁸ The volume sign pattern is very similar to that for volatility, reported in Table 7. The explanatory power is also high, an adjusted R^2 of 0.939. For bubbles, the linear effects have the same signs and slightly greater significance than those reported in the planar regression. The adjusted R^2 rises substantially to 0.903, a clear indication of surface curvature and parameter interaction.

4. Conclusions and future research

Simulation is a standard method for allowing realistic features into an otherwise intractable mathematical model. We use it to investigate how market phenomena are influenced by traders reacting to one another.

To conform with actual markets, we assume that traders cannot fully know the beliefs or resources of competitors; they must attempt to deduce these attributes from trading activity. Beliefs based on private information are revised after observing earlier transactions, particularly transactions executed by agents with substantial resources. Trading progresses sequentially as each agent computes a posterior estimate of value by combining private information with inferences drawn from prior trading about the information possessed by others. The size and direction (buy or sell) of successive transactions are determined by the prevailing price compared with the posterior estimate of value and the trader's resources.

Reacting to fellow traders substantially reduces volatility under many market conditions. On average over all the market conditions (parameters) we studied, price swings were lowered by about 28% in amplitude and return volatility was reduced by approximately 30%. These reductions in volatility are accompanied by a 37.5% reduction in the dollar volume of trading.

A more diffuse prior distribution on value results in a greater reduction of volatility from reactive learning. When the prior is more diffuse, it is weighted less during Bayesian updating; greater weight is placed on the actions of other traders who in aggregate actually do have better information about value.

In contrast, more diversity in private signals reduces the ability of reactive learning to decrease price and return volatility. Traders mislead one another when they are prompted to trade by erroneous private information. In some cases, this can exacerbate the prevalence of bubbles, i.e., excursions by the market price far away from fundamental value. We observe a similar effect when

¹⁸To conserve space, these results will be given to interested readers upon request.

more traders receive no signal at all. They are totally uninformed except by what they learn from watching others, but they trade nonetheless and thereby mislead their competitors.

When traders have limited resources, bias their bids and offers to correct for the winner's curse, or engage in strategically misleading transactions, there is a smaller reduction in volatility induced by learning. (All three of these trading behaviors are subsumed in our trading 'noise' specification.) If such activities are frequently large and other market conditions are not dominant, volatility can actually be increased by learning. In such circumstances, the market would be more stable if traders ignored one another.

When traders react to each other, the market price is a better forecast of value, on average over our parameters by about 12%. Price possesses a lower root mean square prediction error and is closer to value more frequently. The improvement in price accuracy from learning increases with (1) the fraction of traders who receive signals, (2) the accuracy of aggregated private information, and (3) the diversity of private signals.

Notice that greater diversity in private signals is associated with learning increasing volatility and yet improving the accuracy of price as a forecast of value. This interesting result shows that volatility per se is not an appropriate indicator of market quality. When private signals are quite diverse, reacting to others increases volatility but, at the same time, the price established by the end of trading is closer to the asset's true value. Hence, markets with greater volatility can also have more informative prices.

On average, the prevalence of bubbles is reduced by about 13% through learning. Bubble prevalence behaves similarly to volatility, as might have been expected.

Peering deeper into the simulation results, we see that market phenomena such as volatility and price accuracy react in a highly non-linear fashion to market conditions such as heterogeneity of private information and differential resources. In addition, there are strong interaction effects among market conditions. For instance, greater information (signal) diversity reduces the ability of learning to lower volatility, but this effect is mitigated by a more diffuse prior distribution on value. Although many of these findings make intuitive sense a posteriori, they would have been hard to imagine a priori, before seeing the simulation output. A market is indeed a complex, subtle, and delicate contraption.

Several questions for future research are suggested by our results. Among others, we would like to know the answers to the following:

- (a) Would agents be willing to invest to improve the quality of their private signals? If so, how would market quality be affected?
- (b) Are traders who emit misleading signals in early periods able to recoup their losses later?

- (c) How realistic is our assumption of transparency? Would there be benefits in splitting up orders in an effort to hide private information? How would learning be influenced by such behavior?
- (d) How robust are the results to relaxing unrealistic assumptions; e.g., the form of the excess supply function of small traders or the risk neutrality of everyone?
- (e) Would market quality be affected by institutional arrangements such circuit breakers or the imposition of a single-price auction prior to a round of continuous sequential trading?

There are many important and interesting unresolved questions about how traders behave, particularly when they learn from and react to each other. Closed form answers would be very desirable, but may prove a daunting challenge. The numerical approach is much easier for a system as complex as a financial market. Though it cannot provide unequivocal proof of any proposition, simulation may nonetheless offer useful insights and a practical guide for market organizers, traders, and regulators.

References

- Aumann, R.J., 1976. Agreeing to disagree. *The Annals of Statistics* 4 (6), 1236–1239.
- Berger, J.O., 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer, New York.
- De Long, J.B., Schleifer, A., Summers, L.H., Waldmann, R.J., 1990. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance* 45 (2), 379–395.
- Gallant, A. Peter R., Rossi, E., Tauchen, G., 1992. Stock price and volume. *Review of Financial Studies* 5 (2), 199–242.
- Grossman, S., 1976. On the efficiency of competitive stock markets where trades have diverse information. *Journal of Finance* 31 (2), 573–585.
- Hirshleifer, J., 1977. The theory of speculation under alternative regimes of markets. *Journal of Finance* 32, 975–999.
- Karpoff, J.M., 1987. The relation between price changes and trading volume: a survey. *Journal of Financial and Quantitative Analysis* 22 (1), 109–126.
- Keynes, J.M., 1937. *The General Theory of Employment, Interest and Money*. Harcourt, Brace, New York.
- Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. *Journal of Economic Theory* 26, 17–27.
- Scharfstein, D.S., Stein, J.C., 1990. Herd behavior and investment. *American Economic Review* 80 (3), 465–479.
- Shiller, R.J., 1990. Market volatility and investor behavior. *AEA Papers and Proceedings* 80 (2), 58–62.
- Shiller, R.J., 1995. Conversation, information and herd behavior. *AEA Papers and Proceedings* 85 (2), 181–185.
- Tirole, J., 1982. On the possibility of speculation under rational expectations. *Econometrica* 50, 1163–1181.
- Welch, I., 1992. Sequential sales, learning, and cascades. *Journal of Finance* 47 (2), 695–732.
- White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48, 817–838.