

Adaptive Traders and the Design of Financial Markets

SEBASTIEN POUGET*

ABSTRACT

This paper studies a financial market populated by adaptive traders. Learning is modeled following Camerer and Ho (1999). A call market and a Walrasian tatonnement are compared in an environment in which both institutions have the same Nash and competitive equilibrium outcomes. When traders learn via a belief-based model, equilibrium is discovered in both types of markets. In contrast, when traders learn via a reinforcement-based model, convergence to equilibrium is achieved in the Walrasian tatonnement but not in the call market. This paper suggests that market mechanisms can be designed to foster traders' learning of equilibrium strategies.

THE RISK OF TRANSACTING WITH A MORE OR LESS informed investor is inherent to trading in financial markets. This so-called adverse selection risk gives rise to a variety of interesting questions.¹ Do financial market prices reveal investors' private information? Do investors design their trading strategies to mitigate adverse selection risk? Does market organization influence traders' strategies and in turn market efficiency? While these questions have been extensively studied using perfectly rational traders, less attention has been devoted to the case in which traders are boundedly rational.² This paper addresses these issues in the context of financial markets populated by boundedly rational traders.

Starting with Simon (1955), numerous academic studies document the attractiveness of bounded rationality to model decision making.³ One way to incorporate bounded rationality is to represent agents as adaptive organisms, that is,

*University of Toulouse, Institut d'Administration des Entreprises and CRG. I would like to thank Klaus Adam, Vikas Agarwal, Bruno Biais, Stefano Lovo, Lalitha Naveen, Anand Venkateswaran, participants at the 2003 Joint Conference on Information Sciences in Cary, NC, and especially Catherine Casamatta and Naveen Daniel for helpful comments and discussions. I am also grateful to an anonymous referee, an anonymous associate editor, and the Editor, Robert Stambaugh, for their suggestions. Part of this research was done while I was an assistant professor of Finance at the Robinson College of Business, Georgia State University. The financial support of the College is gratefully acknowledged. Anna Agapova and Costanza Meneghetti provided excellent research assistance on part of this paper.

¹ See Grossman (1977) and Hellwig (1980) for seminal contributions to the study of adverse selection risk in financial markets.

² For studies on financial markets with perfectly rational investors, see Grossman and Stiglitz (1980), Glosten and Milgrom (1985), and Kyle (1985) on informational efficiency, Biais, Bossaerts, and Spatt (2003) on trading strategies that mitigate adverse selection, and Madhavan (1992), Pagano and Roell (1996), and Cespa (2004) on market design.

³ See, for example, Conlisk (1996) for a survey on bounded rationality and Camerer (1995) for a survey on individual decision making.

organisms that learn from and make choices according to their past experience. Camerer (2003) provides experimental evidence suggesting that learning models generate more accurate predictions, and hence shed more light on behavior in games than models based on perfect rationality. With respect to financial markets, this research raises questions about price formation in markets populated by adaptive traders. Do adaptive informed traders learn to base their orders on their information? Can adaptive uninformed traders learn to mitigate adverse selection risk? Does market organization affect learning?

To examine these questions, I rely on Camerer and Ho's (1999) Experience Weighted Attraction (EWA) learning model. This model captures the dual nature of adaptive learning, that is, the *law of actual effect* and the *law of simulated effect*. The law of actual effect asserts that chosen actions that are successful will be chosen more often than those that are not successful. This law is at the core of reinforcement learning (see, e.g., Roth and Erev (1995)). The law of simulated effect states that unchosen actions that would have been successful will also be chosen more often. This law is at the core of belief-based learning (see, e.g., Fudenberg and Levine (1998)). Using the EWA model is useful for understanding how the nature of learning affects convergence to equilibrium.

The financial market I consider is in the spirit of Radner (1979). Specifically, the financial market is a pure exchange economy in which agents can trade a risky asset and some traders (the insiders) know the future value of the asset. In order to identify the impact of market design on equilibrium learning, I analyze two market institutions, namely, a call market and a Walrasian tatonnement. These trading architectures are comparable since they both lead to one uniform-price multilateral transaction per trading period. In addition, given the economic environment considered in the present paper, these two institutions have the same fully revealing Perfect Bayesian Equilibrium outcome, and this equilibrium corresponds to the unique competitive rational expectation equilibrium (REE). The idea is to study whether market design influences learning of equilibrium in a context in which it does not affect competitive and strategic equilibrium outcomes.

I study the model described above using computer-based simulations of four treatments that vary according to the type of learning (law of actual effect alone versus laws of actual and simulated effects jointly) and the market mechanism (call market versus Walrasian tatonnement). Each treatment includes 1,000 independent trials and each trial includes 200 successive trading periods or runs. The precise parameters used and the learning process are detailed in Section III.

The results for the call market indicate that if agents rely only on the law of actual effect to adapt their behavior, the system does not converge to the perfectly rational equilibrium benchmark, whereas if agents rely on both the law of actual effect and the law of simulated effect, the market converges toward equilibrium. The rationale is as follows. During the initial runs, insiders have not yet adopted their dominant strategy, which is to reveal their information. This implies that uninformed traders are subject to strategic uncertainty.⁴ As

⁴ For a discussion of strategic uncertainty, see Crawford (1997).

a result, uninformed traders learn *not* to play their equilibrium action. Once insiders have discovered their equilibrium strategy, if agents only reinforce actions that have been successful in the past (law of actual effect only), uninformed traders are unlikely to choose their equilibrium strategy, whereas if agents also use the law of simulated effect, uninformed traders are able to discover the best response to the insiders' behavior.

The results for the Walrasian tatonnement show that the market converges toward equilibrium regardless of whether traders use the law of actual effect alone or whether they use the law of actual effect together with the law of simulated effect. This result obtains for two reasons. First, the Walrasian tatonnement isolates traders from the consequences of strategic uncertainty. Since this market institution is a price-driven mechanism, traders know the potential transaction price before transacting. Uninformed traders can thus learn not to trade at out-of-equilibrium prices. Uninformed traders' equilibrium strategy in the call market is, in contrast, to submit a market order, in which case these agents cannot specify the prices at which they are willing to transact, and hence they are hurt by strategic uncertainty. Second, the Walrasian tatonnement presents information to traders in a simple way: When a potential price is announced, they only need to decide whether they want to trade. This reduces the complexity of the decision problem faced by traders relative to the call market, where traders have to decide whether they want to trade at each potential price.

The present paper is related to studies on pre-trade transparency by Pagano and Roell (1996), Bloomfield and O'Hara (1999), and Flood et al. (1999), among others. It complements this literature by showing that transparency can have beneficial consequences for the ability of boundedly rational traders to design efficient strategies. In particular, this paper suggests that appropriately designed market structure can overcome the consequences of traders' adaptive behavior. This might help explain why stock exchanges that open with a call market, such as the Paris Bourse, the Milan Borsa, the Madrid Bolsa, and the Toronto Stock Exchange, organize a preopening period. This preopening period can help traders determine what the opening transaction price will be, thereby eliminating the price uncertainty inherent to call market mechanisms.

The rest of the paper is organized as follows. Section I discusses the literatures on equilibrium learning and learning in auctions and markets. Section II describes the economic environment and the market structures of interest in this study, along with the perfectly rational theoretical outcomes. Section III presents the learning model and the simulation procedures. The baseline simulation results are given in Section IV, and robustness checks are offered in Section V. Section VI concludes.

I. Related Literature

Gode and Sunder (1993, 1997) show in their seminal contributions that in a setting with private values and zero-intelligence traders, market rules act as a partial substitute for individual rationality. The result in this paper that some market institutions can foster learning is related to this idea. However, in

the present framework institutional features such as the price-driven nature of the Walrasian tatonnement help agents learn equilibrium, whereas in the framework of Gode and Sunder (1993, 1997) agents' behavior does not evolve and market institutions only constrain behavior.

The literature on equilibrium learning consists of two branches. One strand of the literature examines the various types of learning models and their convergence toward equilibrium. For example, Bray (1982), Bray and Savin (1986), and Marcet and Sargent (1989) focus on least-square learning and its theoretical convergence toward REE, Roth and Erev (1995) and Mookherjee and Sopher (1994) study reinforcement learning and its convergence to Nash equilibrium, and Milgrom and Roberts (1990) and Fudenberg and Levine (1998) study belief-based learning and its convergence to Nash equilibrium. The other strand of the literature focuses on mechanism design with bounded rationality. In this second strand, the idea is to design economic institutions that lead to an equilibrium with socially desirable outcomes and then to verify that equilibrium obtains when agents suffer from bounded rationality (see, e.g., Walker (1984), Jordan (1986), and Cabrales (1999)). The present paper contributes to this second strand of the literature by considering financial market design with asymmetric information.

The present paper is also related to the studies of Andreoni and Miller (1995), Arifovic and Ledyard (2003), Routledge (1999), and Feltovich (2003). Andreoni and Miller (1995) simulate the behavior of genetic algorithms in various auction institutions, and Arifovic and Ledyard (2003) simulate adaptive learning in various private-value call markets using evolutionary adaptation. The present paper also investigates the impact of economic institutions on equilibrium discovery. However, it is different from these studies in that it considers a market setting, common values, and individual learning.⁵

Routledge (1999) identifies conditions under which genetic algorithms converge to the REE in a common-value asymmetric information environment. Feltovich (2003) studies agents' learning of equilibrium bids in a common-value auction and finds that when bidders learn using reinforcement, they learn to reduce adverse selection risk very slowly. My results in a market setting are consistent with Feltovich's findings in an auction setting. The present paper also complements these last two studies by showing that the design of market institutions can influence equilibrium discovery. In particular, my result that the sequential nature of the Walrasian tatonnement fosters learning is in line with the conjecture of Ausubel (2004) and the results of Parkes (1999), as the present analysis shows that with asymmetric information, convergence to equilibrium is achieved if market institutions mitigate strategic uncertainty.

The approach used in the present paper is also related to that of Chan et al. (2001), who propose an agent-based model of financial markets with asymmetric information. There are two main differences, however, between Chan et al. (2001) and the present paper. First, I endow traders with various learning

⁵ As noted by Camerer and Ho (1999), the attractiveness of individual learning models is that their parameters have sound psychological underpinnings.

abilities whereas Chan et al. (2001) rely on evolutionary dynamics for the adaptation of trading strategies. Second, I study the impact of market design whereas they focus on the influence of heterogeneous preferences.

The result that uninformed traders' behavior, when they use only the law of actual effect, converges toward an action that protects them from strategic uncertainty echoes the findings of Noe, Rebello, and Wang (2003). They use genetic algorithms to study the design of financing contracts and show that financing contracts that minimize strategic uncertainty tend to emerge and dominate other forms of financing. Consistent with their analysis, in a framework in which the financial contract is given, I conclude that strategic uncertainty is an important determinant of traders' behavior and I study its implications for the design of financial markets.

In addition, my result that Walrasian tatonnement fosters learning is in line with the experimental evidence on sequential auctions (Kagel, Harstad, and Levin (1987), and Kagel and Levin (2001)) and on sequential markets (Pouget (2007)). In this last paper, experimental subjects confronted with similar environments and market institutions are able to discover equilibrium in a Walrasian tatonnement but not in a call market. The present analysis suggests that this might be due to the fact that individuals in experimental financial markets rely mainly on the law of actual effect to adapt their behavior. Studying how various market institutions influence the discovery of optimal strategies is a first step toward the optimal design of financial markets under bounded rationality.

II. Market Settings and Equilibrium with Perfect Rationality

I consider a market for one financial asset. The asset's common value v is a random variable that can take n different values with equal probability. Assume that there are $b_i + s_i$ risk-neutral insiders, where b_i denotes informed buyers and s_i denotes informed sellers, and $b_u + s_u$ risk-neutral uninformed traders, where b_u denotes uninformed buyers and s_u denotes uninformed sellers. Assume that $b_i + s_i = b_u + s_u$. The insiders know the realization of v . In order to generate explicit gains from trade, assume that the asset is worth more for buyers than for sellers. In particular, buyers' valuation for the asset is $v + x$ and sellers' valuation is $v - x$.⁶ Agents have no asset or cash endowment and no budget constraint, and short sales are allowed. There is only one (multilateral) transaction before complete resolution of uncertainty.

Trading can be organized either as a call market or as a Walrasian tatonnement. In the call market, agents submit limit orders to trade one unit of the asset.⁷ The individual orders are combined to compute aggregate supply and demand curves. Transactions occur at the price that maximizes trading volume.

⁶ In order to guarantee that the fully revealing equilibrium price is unique, x is assumed to be strictly smaller than the minimum price variation.

⁷ This reduction in the size of the action space simplifies both the game-theoretical analysis of the game and the burden imposed on the learning process.

If several prices satisfy this condition trading occurs at the price that minimizes aggregate excess demand. If several prices still satisfy this condition, the price is randomly chosen among the admissible potential prices. If there is an order imbalance at the transaction price, random rationing is imposed.⁸

In the Walrasian tatonnement, a first tentative price is announced (equal to the unconditional expectation of the asset's common value). At this price, traders submit market orders to trade one unit of the asset. If the aggregate supply equals the aggregate demand, a transaction occurs. Otherwise, a new tentative price is announced, where the tentative price is increased if demand exceeds supply and the tentative price is decreased if supply exceeds demand. If the market has not cleared after the tenth announcement, the tatonnement process is stopped. Finally, if, after this last price announcement, there is an order imbalance, random rationing is imposed.

When traders are perfectly rational (i.e., maximize their expected utility and formulate rational expectations), the present framework is a particular case of the setting analyzed by Radner (1979). As a result, there exists a unique Rational Expectations Equilibrium (REE). In equilibrium, insiders' information is fully incorporated in prices. The REE is based on the premise that traders are price takers. Since they have only one occasion to trade and profit from their information, informed buyers buy when the price is smaller than or equal to their asset valuation and informed sellers sell when the price is greater than or equal to their asset valuation. This behavior induces prices to incorporate insiders' information. Because insiders' information is perfectly reflected in prices, transacting is individually rational for uninformed traders. In equilibrium in the call market, uninformed buyers submit a quasi-market order to buy (i.e., a buying limit order with a limit price equal to the highest potential asset value), and uninformed sellers submit a quasi-market order to sell. In equilibrium in the Walrasian tatonnement, uninformed agents propose to trade only at prices that correspond to revealing levels.

Following Pouget (2007), one can show that when traders are perfectly rational, the revealing REE is also a Perfect Bayesian Equilibrium. This result holds because traders can only exchange one unit of the asset, which prevents them from having market impact. In the Walrasian tatonnement, this result is also a consequence of the fact that the price adjustment depends only on the sign of the excess demand.

The above discussion implies that, whether one considers price-taking or strategic behavior, a revealing equilibrium exists. Following Pouget (2006), one can show that all the gains from trade are exploited at equilibrium. This equilibrium analysis also implies that because the equilibrium is fully revealing,

⁸ Random rationing proceeds as follows. Consider, for example, the case in which there is one more buyer than the number of sellers at the transaction price. In order for the market to clear, one buyer is randomly chosen (where each buyer has equal probability of being chosen) and excluded from the transaction. The same logic applies when the gap between the number of buyers and sellers is larger, and when there are more sellers than buyers (in which case some sellers are excluded from the transaction).

informed and uninformed traders should earn the same gains from trade (i.e., x per trader).

III. Simulation Setting and Learning Process

A. Simulation Setting⁹

Each simulation contains 1,000 trials. In turn, each trial contains 200 runs. Each run represents one trading session. The common value of the asset is independently drawn from run to run, using the same distribution for all runs. Similarly, each trial represents an independent observation. Agents' behavior, however, is not independent across runs: Agents learn from what happened in previous runs and adapt their behavior accordingly (as is explained below in Section III.B.).

The number n of potential states of the world is set to two, with the common value v equal to three or seven. Potential transaction prices range from one to nine with increments of one. There are eight traders, where $b_i = s_i = b_u = s_u = 2$. The private value x equals one-half.

Given this setting, in the call market, traders have to choose a limit price out of nine potential prices. Traders therefore have nine potential actions. For each action, an attraction and a choice probability are determined by the learning process described below. In the Walrasian tatonnement, a tentative price is announced and then buyers (sellers) have to choose whether to buy (sell) the asset or not. Because there are nine potential prices and at each price an agent can trade or do nothing, traders have 18 potential actions. As indicated above, in each run the first price to be announced is the expected value of the asset, that is, five. Given that a price has been announced, the attraction and the probability of the two potential actions, trading and doing nothing, are determined using the learning process described below.

B. Learning Process

The adaptive learning model considered in this paper is taken from Camerer and Ho (1999). Their so-called EWA model encompasses reinforcement- and belief-based types of learning, that is, this framework captures both the *law of actual effect* and the *law of simulated effect*. The law of actual effect asserts that the likelihood to choose an action is adjusted only if this action has been selected in the previous run: If the action has been successful (unsuccessful), the likelihood that it will be selected in the next run increases (decreases). The law of simulated effect asserts that the likelihood to choose actions that have not been selected in the previous run is adjusted: If the action had been successful

⁹ This setting is chosen to match the experimental design in Pouget (2007), which facilitates comparison between the simulation results, and the behavior observed in the laboratory. One difference between the present setting and Pouget (2007) is that, here, agents are assigned a precise role as buyer or seller. This feature reduces the computational complexity imposed on the learning process.

(unsuccessful) in the previous run, the likelihood that it will be selected in the next run increases (decreases).

The logic of the EWA learning process is that agents make their choices according to the various actions' attractions. For a given agent, these attractions (represented by the function A) correspond to the propensity to choose the various actions. Attractions depend on an agent's experience with the game (represented by the function N) and on past earnings, that is, profits or losses realized or potentially realized by the various actions. More specifically, experience N and attractions A are governed by the two equations,

$$N(t) = \rho \cdot N(t-1) + 1, \quad (1)$$

$$A_i^j(t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(t-1) + [\delta + (1-\delta) \cdot I(a_i^j, a_i(t))] \cdot \pi_i(a_i^j, a_{-i}(t))}{N(t)}. \quad (2)$$

The first equation updates $N(t)$, the measure of an agent's experience with the game after period t has taken place. The parameter ρ depreciates the previous-period experience measure. The second equation updates $A_i^j(t)$, the attraction of each potential action a_i^j , $j = 1, \dots, J$, of player i after period t has taken place. In the present analysis, δ , which controls the nature of learning, is the most important parameter. When $\delta = 0$, agents employ actual effect learning, that is, agent i reinforces only $a_i(t)$, the action chosen in the previous run t , according to the payoff $\pi_i(a_i(t), a_{-i}(t))$ generated by this action, where $a_{-i}(t)$ represents the actions chosen by player i 's opponents in period t . Indeed, the attraction adjustment is driven by the indicator function $I(\cdot, a_i(t))$, which takes the value of one when its two arguments are equal (i.e., when a_i^j , the action to be reinforced, is equal to $a_i(t)$, the action chosen in the previous run), and zero otherwise. When $\delta > 0$, agents also reinforce the actions not chosen in the previous run, in which case δ measures how much the payoffs from these actions influence the attractions. This situation corresponds to simulated effect learning. In particular, when $\delta = 1$, actions' payoffs or potential payoffs equally affect attractions whether or not these actions were chosen in the previous run. The parameter ϕ controls the depreciation of previous attractions. To close the system, the initial values $N(0)$ and $A_i^j(0)$ need to be specified.

Attractions are transformed into choice probabilities according to the following logit model:

$$P_i^j(t+1) = \frac{e^{\lambda \cdot A_i^j(t)}}{\sum_{k=1}^J e^{\lambda \cdot A_i^k(t)}}. \quad (3)$$

The parameter λ represents the sensitivity of players to attractions. These choice probabilities are used to build a cumulative distribution function $F_i(\cdot)$. To determine a trader's actual choice, a random variable u , uniformly drawn between zero and one, is compared to this cumulative distribution function. Action j , $1 < j < J$, is chosen if $F_i(j-1) < u < F_i(j)$; Action 1 is chosen if

$0 < u < F_i(j)$, and action J if $F_i(J - 1) < u < 1$. Camerer and Ho (1999) include a more extensive description of this learning model.

The baseline setting is as follows: δ is set to zero or one depending on the treatment, $\rho = 0$, $\phi = 1$, $\lambda = 1$, $N(0) = 1$, and $A_i^j(0) = 1$. This last parameter choice implies that at the beginning of each trial, each action is equally likely to be chosen. These simulation parameters are consistent with the parameters estimated by Camerer and Ho (1999) using data gathered from several experimental studies. Also, this choice of parameter values has the advantage of simplicity: Attractions represent cumulated past payoffs when only actual effect learning is used (i.e., when $\delta = 0$) and cumulated past potential payoffs when simulated effect learning is also used (i.e., when $\delta = 1$). Indeed, in the baseline setting, the learning model described in equations (1) and (2) boils down to¹⁰

$$A_i^j(t) = A_i^j(t - 1) + [\delta + (1 - \delta) \cdot I(a_i^j, a_i(t))] \cdot \pi_i(a_i^j, a_{-i}(t)). \quad (4)$$

To model their learning process, informed traders are endowed with different sets of attractions $A_i^j(t, v)$ depending on the asset's value. One set, $A_i^j(t, v = 3)$, corresponds to the case in which the value is low, and the other set, $A_i^j(t, v = 7)$, corresponds to the case in which the value is high. When the value is low (high), informed traders use the corresponding set of attractions to randomly determine their action. Despite their adaptive behavior, this allows the informed traders to potentially display different behaviors according to the underlying state of the world. In the simulations, only one set of attractions is updated at the beginning of each run. The attractions that are updated are the ones used in the previous run. Because they do not know the underlying state of the world, uninformed traders have only one set of attractions, and these attractions are updated at the beginning of each round.

In the Walrasian tatonnement, when $\delta > 0$, the learning model updates only the attractions of unchosen actions corresponding to the actual transaction price. For example, if the transaction price is six, the learning model updates the attractions of "trading" or "not trading" at a price of six, but does not update the attractions of "trading" or "not trading" at the other prices.¹¹

To effect the highest level of control on the results, all the randomly generated variables are held constant across simulations with different δ (including the variables used to generate the asset value v and the traders' actions). However, note that the actions themselves will not be the same across trials since the cumulative distribution function $F_i(\cdot)$ evolves endogenously with the learning

¹⁰ To test the sensitivity of the results to the choice of parameter values, I perform various robustness simulations that change one parameter value at a time. The results are reported in Section V below and are similar to those obtained in the baseline case. Simulations with $\delta = 1/2$ are not reported but yield results qualitatively similar to the case in which $\delta = 1$, except that learning is slower.

¹¹ Doing so would require that traders be endowed with elaborate expectations concerning what would have happened if other prices had been reached. Since the focus of this paper is on adaptive learning, such expectations are out of the scope of the present analysis.

process. In addition, the realizations of v and those of the rationing process are identical across simulations with different market structures.

IV. Main Results

A. Market Data

The results for the call market and the Walrasian tatonnement, along with the equilibrium benchmark, are depicted in Figures 1, 2, 3, and 4. Figure 1 graphs informational efficiency, as measured by the absolute difference between the price and the asset value v . Figure 2 plots allocative efficiency, as measured by the gains from trade extracted. Figures 3 and 4 show the gains from trade extracted by informed and uninformed traders, respectively. For each figure, Panels A and B correspond to simulations with $\delta = 0$ and $\delta = 1$, respectively. The variables of interest are averaged across the 1,000 independent trials, giving the evolution of the variables of interest over 200 successive runs. However, in order to limit the amount of data provided in each graph, the figures only present the averages computed over 10 successive runs. For example, the first diamond in Panel A of Figure 1 represents the average absolute difference between the price and the asset value in the call market with $\delta = 0$ computed using runs 1 through 10, the second diamond represents the average absolute difference computed using runs 11 through 20, etc. Visual inspection of these figures suggests that (i) simulated effect learning is necessary for traders to discover the REE in the call market, whereas (ii) with actual effect learning only, traders are able to learn and play the REE in the Walrasian tatonnement.

To confirm these graphical evidence, Tables I through IV present results of statistical tests related to price efficiency, allocative efficiency, and gains from trade extracted by informed and by uninformed traders, respectively. These tables only include every 20th run. Panels A present the variable of interest averaged across 1,000 trials. Standard deviations across trials are also provided. In these panels, there are two comparisons of interest. First, one can look at the difference between what happens when $\delta = 0$ and when $\delta = 1$ both in the call market and in the Walrasian tatonnement. This enables one to compare learning when traders use only the law of actual effect to learning when traders use both the law of actual effect and the law of simulated effect. Second, one can look at the difference between what happens in the call market and in the Walrasian tatonnement under the various learning models. This sheds some light on the role of market design when traders exhibit adaptive behavior. These comparisons are performed using t -tests that are reported in Panels B of Tables I through IV.

In the call market, when traders use only the law of actual effects ($\delta = 0$), the market does not converge to the REE, as indicated by the fact that the informational efficiency, gains from trade, and profits of uninformed traders are far from the REE levels (see the column corresponding to the call market with $\delta = 0$ in Panel A of Tables I, II, and IV). In addition, Panel A of Table III shows





Figure 1. Price efficiency in the baseline simulations. Price efficiency is measured by the absolute deviation between the price and the common value of the asset. For each run, this absolute deviation is averaged across 1,000 trials. Figure 1 only presents averages computed over 10 successive runs. Gray diamonds correspond to the call market and black squares to the Walrasian tatonnement. The equilibrium benchmark is represented by circles.

that insiders' profits decrease over time. This is because uninformed traders are learning to avoid loss-making actions instead of playing according to the equilibrium. This behavior reduces the liquidity of the market and prevents insiders from extracting all the gains from trade. When traders use the law of actual effect and the law of simulated effect ($\delta = 1$), the market discovers




Figure 2. Allocative efficiency in the baseline simulations. Allocative Efficiency is measured by the gains from trade extracted. For each run, these gains from trade are averaged across 1,000 trials. Figure 2 only presents averages computed over 10 successive runs. Gray diamonds correspond to the call market and black squares to the Walrasian tatonnement. The equilibrium benchmark is represented by circles.

equilibrium. This is indicated by the fact that the average informational efficiency, gains from trade, and profits of both informed and uninformed traders converge to the REE (and end up being not statistically significantly different from this level; see the column corresponding to the call market with $\delta = 1$ in Panel A of Tables I, II, III, and IV). These results are confirmed by the first




Figure 3. Informed traders' gains from trade in the baseline simulations. For each run, informed gains from trade are averaged across 1,000 trials. Figure 3 only presents averages computed over 10 successive runs. Gray diamonds correspond to the call market and black squares to the Walrasian tatonnement. The equilibrium benchmark is represented by circles.

column of Panel B in Tables I through IV, which shows that the variables of interest are significantly closer to the REE when $\delta = 1$ than when $\delta = 0$.¹²

¹² Note that for Table I, a positive difference between the cases $\delta = 0$ and $\delta = 1$ indicates that when $\delta = 0$, the mean absolute deviation is higher (and thus the market is further away from the equilibrium) than when $\delta = 1$. The opposite reasoning is true for Tables II, III, and IV.




Figure 4. Uninformed traders' gains from trade in the baseline simulations. For each run, uninformed gains from trade are averaged across 1,000 trials. Figure 4 only presents averages computed over 10 successive runs. Gray diamonds correspond to the call market and black squares to the Walrasian tatonnement. The equilibrium benchmark is represented by circles.

In the Walrasian tatonnement, Tables I through IV suggest that, whether traders use only the law of actual effect or both the law of actual effect and the law of simulated effect, the market converges to the REE. This is illustrated in the two columns of Panels A that correspond to the Walrasian tatonnement with $\delta = 0$ and with $\delta = 1$. Indeed, these columns indicate that the average informational efficiency, gains from trade, and profits of both informed

Table I
Informational Efficiency of Prices

Informational efficiency is measured by the absolute deviation between prices and the common value of the asset. This table only includes every 20th run. Panel A presents the absolute deviation averaged across 1,000 trials. Standard deviations across trials are also provided. The means that are *not* statistically significantly different from the equilibrium level are indicated by an asterisk. Panel B presents the *t*-statistics for difference in means tests applied to data in Panel A. CM and WT stand for call market and Walrasian tatonnement, respectively.

Panel A: Mean Absolute Deviation between Price and Value									
Runs	Call Market				Walrasian Tatonnement				Equilibrium
	<i>delta</i> = 0		<i>delta</i> = 1		<i>delta</i> = 0		<i>delta</i> = 1		
	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	
20	1.64	1.09	0.52	0.86	0.81	0.84	0.41	0.70	0.00
40	1.26	0.95	0.09	0.35	0.55	0.70	0.23	0.60	0.00
60	1.04	0.92	0.01*	0.19	0.26	0.54	0.08	0.28	0.00
80	0.80	0.77	0.00*	0.00	0.15	0.43	0.08	0.37	0.00
100	1.04	1.09	0.01	0.07	0.17	0.47	0.01	0.08	0.00
120	0.72	0.77	0.00*	0.00	0.07	0.31	0.02	0.17	0.00
140	0.72	0.77	0.00*	0.00	0.09	0.34	0.00*	0.03	0.00
160	0.83	0.96	0.00*	0.00	0.06	0.30	0.00*	0.06	0.00
180	0.93	1.11	0.00*	0.07	0.06	0.30	0.00*	0.00	0.00
200	1.03	1.24	0.00*	0.00	0.01	0.15	0.00*	0.07	0.00

Panel B: <i>t</i> -statistics				
Runs	Call Market <i>delta</i> = 0 vs. <i>delta</i> = 1	Walrasian Tatonnement <i>delta</i> = 0 vs. <i>delta</i> = 1	<i>delta</i> = 0 CM vs. WT	<i>delta</i> = 1 CM vs. WT
20	25.44	11.68	18.97	3.16
40	36.57	10.89	19.15	−6.42
60	34.53	9.52	23.05	−6.08
80	32.80	3.80	23.15	−7.16
100	30.22	10.59	23.45	−0.58
120	29.58	4.72	24.77	−2.93
140	29.75	7.70	23.98	−1.00
160	27.28	5.90	24.13	−1.00
180	26.42	5.94	24.10	1.34
200	26.31	2.05	25.75	−1.34

and uninformed traders converge toward the REE under both specifications of the learning model.¹³ However, it appears that when $\delta = 1$, the convergence is slightly stronger. This is suggested by the fact that (i) when $\delta = 1$, the variables of interest eventually become statistically indistinguishable from the REE

¹³ This result is not surprising since, in the Walrasian tatonnement, only one action other than the chosen action gets updated. I argue below that the low number of potential actions to be updated in this market structure is one reason why it converges toward equilibrium whether traders use actual effect learning alone or together with simulated effect learning.

Table II
Allocative Efficiency

Gains from trade extracted are measured by the sum of traders' profits. This table only includes every 20th run. Panel A presents the mean and standard deviation across 1,000 trials of the gains from trade extracted. The means that are *not* significantly different from the equilibrium level are indicated by an asterisk. Panel B presents the *t*-statistics for difference in means tests applied to data in Panel A. CM and WT stand for call market and Walrasian tatonnement, respectively.

Panel A: Gains from Trade Extracted									
Runs	Call Market				Walrasian Tatonnement				Equilibrium
	$\delta = 0$		$\delta = 1$		$\delta = 0$		$\delta = 1$		
	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	
20	2.32	0.70	3.09	0.63	2.22	0.65	2.50	1.05	4.00
40	2.26	0.88	3.83	0.48	2.78	1.00	3.51	1.04	4.00
60	2.00	0.79	3.98	0.18	3.38	1.00	3.74	0.92	4.00
80	1.96	0.81	4.00*	0.04	3.61	0.90	3.91	0.44	4.00
100	1.96	0.85	4.00*	0.00	3.64	0.90	3.99	0.17	4.00
120	1.95	0.80	4.00*	0.00	3.81	0.70	3.98	0.21	4.00
140	1.85	0.81	4.00*	0.00	3.78	0.81	4.00*	0.06	4.00
160	1.86	0.81	4.00*	0.00	3.86	0.58	4.00*	0.03	4.00
180	1.78	0.90	4.00*	0.00	3.90	0.49	4.00*	0.00	4.00
200	1.77	0.84	4.00*	0.00	3.97	0.26	4.00*	0.07	4.00

Panel B: t -statistics				
Runs	Call Market $\delta = 0$ vs. $\delta = 1$	Walrasian Tatonnement $\delta = 0$ vs. $\delta = 1$	$\delta = 0$ CM vs. WT	$\delta = 1$ CM vs. WT
20	−25.92	−7.05	3.32	15.28
40	−49.56	−16.08	−12.35	8.69
60	−77.28	−8.39	−34.19	7.97
80	−79.63	−9.22	−43.16	6.49
100	−76.04	−12.07	−42.81	2.42
120	−80.97	−7.49	−55.24	2.60
140	−83.94	−8.64	−53.15	1.00
160	−83.35	−7.36	−63.54	1.00
180	−78.10	−6.64	−65.57	0.00
200	−83.47	−3.66	−78.59	1.34

levels, and (ii) as underscored by the second column of Panels B, the market is generally significantly closer to the REE when $\delta = 1$ than when $\delta = 0$.

The results reported above suggest that the Walrasian tatonnement fosters convergence toward the REE. To further establish this point, the last two columns of Panel B in Tables I through IV present *t*-tests of the difference in means of the variables of interest when trading occurs in a call market and when it occurs in a Walrasian tatonnement. These tests suggest that when traders use only the law of actual effect, appropriately choosing the market structure is crucial to ensure convergence toward the REE. Indeed, the third column of Panel B indicates that the average informational efficiency, gains from trade, and profits of both informed and uninformed traders are significantly closer

Table III
Surplus of Informed Traders

The surplus of informed traders is measured by the sum of informed traders' profits. This table only includes every 20th run. Panel A presents the mean and standard deviation across 1,000 trials of the gains from trade extracted by informed traders. The means that are *not* significantly different from the equilibrium level are indicated by an asterisk. Panel B presents the *t*-statistics for difference in means tests applied to data in Panel A. CM and WT stand for call market and Walrasian tatonnement, respectively.

Panel A: Gains from Trade Extracted by Informed Traders									
Runs	Call Market				Walrasian Tatonnement				Equilibrium
	$\delta = 0$		$\delta = 1$		$\delta = 0$		$\delta = 1$		
	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	
20	1.81	1.58	1.85	1.08	1.77	1.03	1.54	0.90	2.00
40	1.55	1.25	1.97	0.25	1.77	1.17	1.99*	0.90	2.00
60	1.53	1.20	2.00*	0.25	1.93	0.67	1.89	0.48	2.00
80	1.37	1.08	2.00*	0.00	1.96	0.59	2.07	0.46	2.00
100	1.48	1.22	2.00*	0.00	1.94	0.53	2.00*	0.09	2.00
120	1.35	1.06	2.00*	0.00	1.95	0.41	2.01*	0.15	2.00
140	1.29	1.10	2.00*	0.00	1.93	0.45	2.00*	0.03	2.00
160	1.38	1.18	2.00*	0.00	1.98*	0.37	2.00*	0.05	2.00
180	1.34	1.21	2.00*	0.00	2.00*	0.27	2.00*	0.00	2.00
200	1.53	1.32	2.00*	0.00	2.00*	0.20	2.00*	0.10	2.00

Panel B: t -statistics				
Runs	Call Market $\delta = 0$ vs. $\delta = 1$	Walrasian Tatonnement $\delta = 0$ vs. $\delta = 1$	$\delta = 0$ CM vs. WT	$\delta = 1$ CM vs. WT
20	−0.68	5.42	0.62	7.03
40	−10.33	−4.89	−3.92	−0.78
60	−12.26	1.52	−9.28	6.53
80	−18.43	−4.38	−15.20	−4.52
100	−13.57	−3.32	−11.12	0.19
120	−19.59	−3.98	−16.90	−1.07
140	−20.53	−4.81	−17.23	−1.00
160	−16.70	−1.49	−15.54	−1.00
180	−17.17	0.47	−16.87	0.00
200	−11.12	−0.78	−10.95	−1.15

to equilibrium in the Walrasian tatonnement than in the call market. When traders use both the law of actual effect and the law of simulated effect, however, market structure appears not to have a clear impact on equilibrium discovery. This result is not surprising in light of the fact that, under both market architectures, the system converges toward equilibrium.

B. Individual Traders' Choice Probabilities

To understand the forces at work behind the results presented above, it is instructive to analyze agents' choices of action. Consider first the case in which

Table IV
Surplus of Uninformed Traders

The surplus of uninformed traders is measured by the sum of uninformed traders' profits. This table only includes every 20th run. Panel A presents the mean and standard deviation across 1,000 trials of the gains from trade extracted by uninformed traders. The means that are *not* significantly different from the equilibrium level are indicated by an asterisk. Panel B presents the *t*-statistics for difference in means tests applied to data in Panel A. CM and WT stand for call market and Walrasian tatonnement, respectively.

Panel A: Gains from Trade Extracted by Uninformed Traders									
Runs	Call Market				Walrasian Tatonnement				Equilibrium
	<i>delta</i> = 0		<i>delta</i> = 1		<i>delta</i> = 0		<i>delta</i> = 1		
	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	Av.	St. Dev.	
20	0.51	1.47	1.24	1.26	0.44	1.01	0.96	0.93	2.00
40	0.70	1.29	1.86	0.47	1.01	1.39	1.52	1.24	2.00
60	0.47	1.18	1.98	0.33	1.45	1.03	1.85	0.51	2.00
80	0.59	1.04	2.00*	0.04	1.65	0.91	1.84	0.76	2.00
100	0.48	1.25	2.00*	0.00	1.69	0.81	1.99	0.15	2.00
120	0.61	1.07	2.00*	0.00	1.86	0.56	1.98	0.23	2.00
140	0.56	1.09	2.00*	0.00	1.84	0.61	2.00*	0.09	2.00
160	0.48	1.23	2.00*	0.00	1.88	0.58	2.00*	0.08	2.00
180	0.43	1.24	2.00*	0.00	1.89	0.52	2.00*	0.00	2.00
200	0.24	1.37	2.00*	0.00	1.97	0.32	1.99*	0.17	2.00

Panel B: <i>t</i> -statistics				
Runs	Call Market		Walrasian Tatonnement	
	<i>delta</i> = 0 vs. <i>delta</i> = 1		<i>delta</i> = 0 vs. <i>delta</i> = 1	
			<i>delta</i> = 0	<i>delta</i> = 1
			CM vs. WT	CM vs. WT
20		−11.86	−11.76	1.12
40		−26.69	−8.64	−5.17
60		−38.71	−11.03	−19.68
80		−42.79	−5.05	−24.26
100		−38.55	−11.28	−25.75
120		−41.11	−6.22	−32.77
140		−41.63	−7.90	−32.38
160		−38.99	−6.38	−32.50
180		−40.09	−6.39	−34.41
200		−40.73	−2.25	−38.96

trading is organized as a call market. To interpret the results, it is useful to detail the perfectly rational equilibrium actions. In equilibrium with perfectly rational agents, informed buyers (sellers) submit a limit order to buy (sell) at a price lower (higher) than or equal to the asset value, that is, three or seven. Perfectly rational uninformed buyers (sellers) anticipate this and submit a quasi-market order to buy (sell).¹⁴

¹⁴ Uninformed buyers (sellers) have to submit a limit price lower (higher) than or equal to the highest (lowest) potential value seven (three). This is what is referred to as a quasi-market order.

The results presented above suggest that the law of actual effect ($\delta = 0$) is not sufficient to induce adaptive traders to discover equilibrium. Figure 5 sheds some light on this issue. This figure displays the evolution of choice probabilities for four representative agents, namely, an informed buyer (Panel A), an informed seller (Panel B), an uninformed buyer (Panel C), and an uninformed seller (Panel D).¹⁵ Figure 5 indicates that during the initial runs, insiders have not yet adopted their dominant strategy, which is to reveal their information (see the first 100 runs in Figure 5, Panel A when $v = 3$ and Panel B when $v = 7$). This implies that uninformed traders are subject to strategic uncertainty (see Crawford (1997)).

In the present framework, it turns out that the uninformed traders' equilibrium action is not a dominant strategy; indeed, it is strictly dominated. Consistent with this idea, during the initial runs uninformed traders learn not to play their equilibrium action, but instead, to play an action that protects them from strategic uncertainty (see the first 100 runs in Figure 5, Panels C and D). For an uninformed buyer, this action is to buy up to a price of three (which is the minimum value of the asset) and not above. For an uninformed seller, this action is to sell up to a price of seven (which is the maximum value of the asset) and not below.

Consistent with the fact that their equilibrium strategy is a dominant strategy, insiders eventually learn to play their equilibrium action (see the last 100 runs in Figure 5, Panel A when $v = 3$ and Panel B when $v = 7$).¹⁶ However, since agents only reinforce actions that have been successful in the past, uninformed traders' unchosen equilibrium strategy remains unlikely to be chosen. In line with this interpretation, the results show that by run 200, the action that protects uninformed traders from strategic uncertainty is chosen with a probability slightly higher than 25% while their equilibrium action is chosen with a probability of 13%.

In summary, what prevents uninformed traders from discovering their equilibrium strategy in the call market is that, in equilibrium, they have to submit a quasi-market order. In a call market, the strategic uncertainty induced by insiders not choosing their equilibrium strategy translates into price uncertainty. Given that the price will not always be set at a revealing level, it is no more advantageous for uninformed traders to submit a quasi-market order. Since, in early runs, quasi-market orders if submitted are harmful, uninformed traders will learn not to choose them. Once insiders have learned to play their equilibrium strategy (i.e., when they reveal their information), this initial learning pattern prevents uninformed traders from discovering their equilibrium strategy.

¹⁵ The agents not displayed in Figure 5 exhibit the same dynamics. This is also the case for Figures 6, 7, and 8.

¹⁶ Figure 5, Panel A when $v = 7$ and Panel B when $v = 3$ suggest that informed traders' learning dynamics lead them to try and in turn benefit from transacting at disequilibrium prices. This does not disturb price efficiency (see Figure 1) because, as indicated by Figure 5, at least one type of informed traders (the buyers or the sellers) will push prices toward the fundamental value.




Figure 5. Evolution of choice probabilities in the call market when $\delta = 0$ in the baseline simulations. For each run, the data are averaged across 1,000 trials. Figure 5 only presents averages computed over 20 successive runs. Agents can choose one out of nine orders with limit price ranging from one to nine. For informed agents, Figure 5 plots choice probabilities conditional on whether the value of the asset is three or seven. The data correspond to a representative agent in each category. Equilibrium actions are indicated in dark gray.

When agents rely on both the law of actual effect and the law of simulated effect ($\delta = 1$), for the same parameterizations as earlier, the market converges toward equilibrium. This phenomenon is illustrated in Figure 6. In this case, insiders quickly learn to play their equilibrium (dominant strategy) action, which is to reveal their information (see Figure 6, Panels A and B). Uninformed traders using the law of simulated effect are then able to quickly discover the best response to insiders' behavior (see Figure 6, Panels C and D).


Turning to the case in which trading is organized as a Walrasian tatonnement, the market converges to equilibrium both when agents use only the law of actual effect and when they use the laws of actual and simulated effects together. This result, illustrated in Figures 7 and 8, is obtained for two reasons. First, in a Walrasian tatonnement, potential transaction prices are announced before transactions occur. This implies that there is no price uncertainty and that agents are less subject to strategic uncertainty. As indicated by Figure 7 (Panels C and D), even when agents use only the law of actual effects, uninformed agents discover their equilibrium strategies because they can specify the prices at which they are willing to trade (i.e., the fully revealing price levels three and seven). It is clear from these graphs that uninformed agents are able to avoid transacting at intermediate prices (i.e., at prices of four, five, and six). When agents use both the laws of actual and simulated effects, Figure 8 shows the same phenomenon with a slight variation, namely, uninformed traders can afford to transact at some intermediate prices because they learn that when prices reach these levels, it is profitable for them to trade.¹⁷ Second, in the Walrasian tatonnement, information is presented to traders in a simple way: When a potential price is announced, they only need to decide if they want to trade or not. This reduces the complexity of the task faced by traders relative to the call market, where traders have to decide whether they want to trade at each potential price. The burden imposed on the learning process is thus reduced, which enables traders to discover the appropriate behavior quickly.

V. Robustness Tests


This section first discusses simulations that are intended to check the technical robustness of the results. Additional simulations that are designed to test the validity of the interpretations are then presented. In all these analyses, only one feature is changed at a time.¹⁸

¹⁷ In Figures 7 and 8, some choice probabilities may appear anomalous. For example, in Figure 7, Panel A, when $v = 3$, the probability of informed buyers buying is around 50% if the announced price is nine; such an action would clearly be unprofitable and the informed buyers should learn to avoid it. The anomalous probabilities reflect the fact that, in the Walrasian tatonnement, the learning model does not update actions corresponding to a price that was not the transaction price. Because some prices are almost never reached, the agents do not learn how to behave in these circumstances. This issue of visiting different states in the learning process is related to the more general problems of sufficient exploration in Q-learning (see, for example, Bertsekas and Tsitsiklis (1996)).


¹⁸ Some of the results discussed in this section are not reported but are available upon request.



Panel A: Informed Buyer



Panel B: Informed Seller



Panel C: Uninformed Buyer

Panel D: Uninformed Seller

Figure 6. Evolution of choice probabilities in the call market when $\delta = 1$ in the baseline simulations. For each run, the data are averaged across 1,000 trials. Figure 6 only presents averages computed over 20 successive runs. Agents can choose one out of nine orders with limit price ranging from one to nine. For informed agents, Figure 6 plots choice probabilities conditional on whether the value of the asset is three or seven. The data correspond to a representative agent in each category. Equilibrium actions are indicated in dark gray.




Figure 7. Evolution of choice probabilities in the Walrasian tatonnement when $\delta = 0$ in the baseline simulations. For each run, the data are averaged across 1,000 trials. Figure 7 only presents averages computed over 20 successive runs. Figure 7 indicates the evolution of the probability that agents choose to trade conditional on one of the nine potential prices being announced. The complementary probability of choosing the alternative action, that is, doing nothing, is not shown in Figure 7. For informed agents, Figure 7 plots the choice probabilities conditional on whether the value of the asset is three or seven. The data correspond to a representative agent in each category. Probabilities that are equal to one in equilibrium are indicated in dark gray (other probabilities are equal to zero in equilibrium).




Figure 8. Evolution of choice probabilities in the Walrasian tatonnement when $\delta = 1$ in the baseline simulations. For each run, the data are averaged across 1,000 trials. Figure 8 only presents averages computed over 20 successive runs. Figure 8 indicates the evolution of the probability that agents choose to trade conditional on one of the nine potential prices being announced. The complementary probability of choosing the alternative action, that is, doing nothing, is not shown in Figure 8. For informed agents, Figure 8 plots the choice probabilities conditional on whether the value of the asset is three or seven. The data correspond to a representative agent in each category. Probabilities that are equal to one in equilibrium are indicated in dark gray (other probabilities are equal to zero in equilibrium).

A. Technical Robustness

First, I modify the baseline environment to consider a situation in which there are three potential asset values instead of two. In these simulations, the asset can be worth three, five, or seven with equal probabilities. The objective is to determine whether additional complexity alters the main results. Second, the asset values of the baseline environment are modified. In these simulations, the asset can be worth two or eight with equal probabilities instead of three and seven. This test checks whether the baseline results hold when asset values are more extreme. Third, random gains from trade are introduced in the baseline environment. Instead of being constant and equal to one-half, traders' private value x is independently randomly drawn across traders at the beginning of each run using a uniform distribution with realizations between zero and one. The objective is to check whether this additional noise prevents convergence to equilibrium. Finally, I perform a series of simulations that change one parameter value at a time using the following values: $\lambda \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1.2, 1.4, 1.6, 1.8, 2\}$, $\Phi \in \{0.80, 0.85, 0.90, 0.95\}$, and $\rho \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$.¹⁹ All these simulations yield results that are qualitatively similar to the baseline results. Thus, unlike other learning models such as Brock and Hommes (1997), the results in this paper are not very sensitive to the learning parameters.

B. Uninformed Traders Start Close to Equilibrium

The objective of this subsection is to gauge the severity of the problem posed by strategic uncertainty when traders use actual effect learning. I modify the baseline environment to consider a situation in which uninformed agents' initial choice probabilities are close to their equilibrium strategy. As in the baseline case, insiders are not initially programmed to play equilibrium. Uninformed traders are thus not isolated from strategic uncertainty.

Attention is restricted to the treatment corresponding to the call market with $\delta = 0$.²⁰ The initial attraction of uninformed agents' equilibrium action is set to four while other actions' attractions are left unchanged at one. Consider, for example, an uninformed buyer. Her initial attractions are: four for the order with a limit price of seven, and one for the remaining eight orders. Applying the logit transformation presented in equation (3) with $\lambda = 1$ yields choice probabilities of 71.52% for the equilibrium order and 3.56% for each of the remaining eight orders. Clearly, this uninformed trader is initially likely to choose her equilibrium action.²¹

Panel A of Figure 9 uses the format of Figure 5, Panel C, and presents the evolution of choice probabilities for a representative uninformed buyer (the results for the other uninformed traders are similar). It appears that, over

¹⁹ Simulations that vary the initial attraction $A_i^j(0)$ and experience measure $N(0)$ are not performed since it can be shown that in the present framework they do not affect choice probabilities.

²⁰ For the other three treatments, convergence toward equilibrium obtains.

²¹ In the call market, in equilibrium with perfect rationality, uninformed buyers assign a probability of one to the order with a limit price of seven and a null probability to the other orders.




Figure 9. Evolution of uninformed buyers' choice probabilities when $\delta = 0$ in the call market in various robustness simulations. For each run, the data are averaged across 1,000 trials. Figure 9 only presents averages computed over 20 successive runs. Figure 9 displays the evolution of choice probabilities for a representative uninformed buyer when $\delta = 0$. Equilibrium actions are indicated in dark gray. Panel A corresponds to the case in which uninformed traders start close to equilibrium; Panel B to the case in which all traders except the representative uninformed buyer play equilibrium.

time, the probability of choosing the equilibrium action decreases while the probability of choosing the order to buy up to a price of three increases. This is in line with the result that uninformed agents learn not to play equilibrium during the initial runs, that is, when insiders do not reveal their information. This result also confirms the later finding that uninformed traders learn to protect themselves from strategic uncertainty. Overall, this subsection shows that strategic uncertainty impedes equilibrium discovery in the call market since, even when they start close to equilibrium, uninformed traders are not able to learn equilibrium.

C. All Traders Play Equilibrium Except One Uninformed Agent

This subsection studies a case in which there is no strategic uncertainty. In particular, I modify the baseline environment to consider a situation in which all traders play equilibrium except for one uninformed buyer who learns according to the EWA model specified in equation (3).²² This uninformed buyer therefore faces no strategic uncertainty.

Attention is restricted to the treatment corresponding to the call market with $\delta = 0$.²³ Panel B of Figure 9 presents the evolution of choice probabilities for the uninformed buyer who faces no strategic uncertainty. This figure shows that this uninformed buyer is able to discover the equilibrium strategy. Indeed,

²² The results are similar for the case in which all the traders play equilibrium except for one uninformed seller.

²³ For the other three treatments, convergence toward equilibrium obtains.

starting from the 40th run, the three orders that offer to buy at prices of seven, eight, or nine are chosen around 90% of the time. Note that when all the other traders play equilibrium, buying up to a price of eight or nine yields the same payoff as buying up to a price of seven. This is because the uninformed buyer's action does not affect the transaction price, which equals three or seven depending on the asset value. This result confirms that the reason uninformed traders do not discover equilibrium in the baseline simulation in the call market is that they are learning to cope with the noise created by other traders not playing their equilibrium strategies. This result also suggests that the Walrasian tatonnement is able to effectively protect traders from strategic uncertainty.

VI. Conclusion

This paper studies the role of market institutions in the discovery of equilibrium when traders adaptively adjust their strategies. In a Radner (1979) type financial market with asymmetric information, I simulate traders' behavior using Camerer and Ho's (1999) EWA learning model, and I compare the performance of a call market and a Walrasian tatonnement. The environment is such that both institutions have the same (competitive and game-theoretical) equilibrium outcomes. This framework is used to show how market design affects learning in a context in which it does not affect equilibrium. The main result of the paper suggests that when traders reinforce only the actions that have successfully been chosen in the past, two characteristics of market design are jointly sufficient to ensure convergence to equilibrium: Market institutions should mitigate strategic uncertainty and be computationally simple. These characteristics are present in the Walrasian tatonnement but not in the call market, where traders submit limit orders.

Could a call market replicate the advantages of the Walrasian tatonnement? A promising direction for call market design is to allow traders to post stop orders in addition to limit orders, as the use of stop orders could allow traders to eliminate the risk of transacting at out-of-equilibrium prices.^{24,25} However, one problem with this design is that introducing stop orders significantly increases the computational complexity faced by traders. This design thus captures an important aspect of the Walrasian tatonnement (the reduction of strategic uncertainty), but at the expense of another important characteristic (computational simplicity). A more extensive investigation of call market design is left for future research.

It would also be interesting to generalize the approach proposed in the present paper to other market mechanisms. An interesting avenue for future work

²⁴ A buying (selling) stop order is an order to buy (sell) at prices at and above (below) a particular threshold.

²⁵ In the present paper, uninformed buyers' equilibrium strategy is less exposed to strategic uncertainty if they can offer to buy one unit of the asset at prices of three and seven but not at other prices. This demand schedule can be implemented with the following menu of orders: a limit order to buy at or below three, a stop order to sell at or below two, a limit order to buy at or below seven, and a stop order to sell at or below six (all these orders being valid for one unit).

would be to examine the performance of dealer markets to understand how features such as centralization and transparency affect the ability of adaptive traders to discover equilibrium. One could also study the comparative ability of dealer and auction markets to induce equilibrium learning. Overall, this line of research would contribute to the optimal design of financial markets populated by boundedly rational traders.

REFERENCES

- Andreoni, James, and John H. Miller, 1995, Auctions with artificial adaptive agents, *Games and Economic Behavior* 10, 39–64.
- Arifovic, Jasmina, and John Ledyard, 2003, Information and dynamics: Sequences of call markets, *Information Systems Frontiers* 5, No. 1, 39–45.
- Ausubel, Lawrence M., 2004, An efficient ascending-bid auction for multiple objects, *American Economic Review* 94, No. 5, 1452–1475.
- Bertsekas, Dimitri P., and John N. Tsitsiklis, 1996, *Neuro-Dynamic Programming* (Nashua, NH: Athena Scientific).
- Biais, Bruno, Peter Bossaerts, and Chester Spatt, 2003, Equilibrium asset pricing under heterogeneous information, IDEI Working Paper, No. 159.
- Bloomfield, Robert, and Maureen O'Hara, 1999, Market transparency: Who wins and who loses? *The Review of Financial Studies* 12, 5–35.
- Bray, Margaret, 1982, Learning, estimation, and the stability of rational expectations, *Journal of Economic Theory* 26, 318–339.
- Bray, Margaret, and N. Eugene Savin, 1986, Rational expectations equilibria, learning, and model specification, *Econometrica* 54, 1129–1160.
- Brock, William A., and Cars H. Hommes, 1997, A rational route to randomness, *Econometrica* 65, 1059–1097.
- Cabrales, Antonio, 1999, Adaptive dynamics and the implementation problem with complete information, *Journal of Economic Theory* 86, 159–184.
- Camerer, Colin F., 1995, Individual decision making, in John H. Kagel and Alvin E. Roth, eds.: *The Handbook of Experimental Economics* (Princeton, NJ: Princeton University Press), 587–703.
- Camerer, Colin F., 2003, *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press).
- Camerer, Colin F., and Teck-Hua Ho, 1999, Experience-weighted attraction learning in normal form games, *Econometrica* 67, 827–874.
- Cespa, Giovanni, 2004, A comparison of stock market mechanisms, *RAND Journal of Economics* 35, 803–824.
- Chan, Nicholas, Blake LeBaron, Andrew W. Lo, and Tomaso Poggio, 2001, Agent-based models of financial markets: A comparison with experimental markets, Working paper No. 4195-01, MIT Sloan School of Management.
- Conlisk, John, 1996, Why bounded rationality? *Journal of Economic Literature* 34, 669–700.
- Crawford, Vincent P., 1997, Theory and experiment in the analysis of strategic interaction, in David Kreps and Ken Wallis, eds.: *Advances in Economics and Econometrics: Theory and Applications I* (New York: Cambridge University Press), 206–242.
- Feltovich, Nicholas J., 2003, Slow learning in the market for lemons: A note on reinforcement learning and the winner's curse, *Proceedings of the 7th Joint Conference on Information Sciences*, 1023–1026.
- Flood, Mark D., Ronald Huisman, Kees G. Koedijk, and Ronald J. Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *Review of Financial Studies* 12, 37–59.
- Fudenberg, Drew, and David K. Levine, 1998, *Theory of Learning in Games* (MIT Press, MA).
- Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.

- Gode, Dhananjay K., and Shyam Sunder, 1993, Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy* 101, 119–137.
- Gode, Dhananjay K., and Shyam Sunder, 1997, What makes markets allocationally efficient? *Quarterly Journal of Economics* 112, In Memory of Amos Tversky (1937–1996), 603–630.
- Grossman, Sanford J., 1977, The existence of futures markets, noisy rational expectations and informational externalities, *Review of Economic Studies* 44, 431–449.
- Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.
- Hellwig, Martin F., 1980, On the aggregation of information in competitive markets, *Journal of Economic Theory* 22, 477–498.
- Jordan, James, 1986, Instability in the implementation of Walrasian allocations, *Journal of Economic Theory* 39, 301–328.
- Kagel, John H., Ronald M. Harstad, and Dan Levin, 1987, Information impact and allocation rules in auctions with affiliated private values: A laboratory study, *Econometrica* 55, 1275–1304.
- Kagel, John H., and Dan Levin, 2001, Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickrey auctions, *Econometrica* 69, 1275–1304.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- Madhavan, Ananth, 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607–641.
- Marcet, Albert, and Thomas J. Sargent, 1989, Convergence of least squares learning mechanisms in self referential linear stochastic models, *Journal of Economic Theory* 48, 337–368.
- Milgrom, Paul, and John Roberts, 1990, Rationalizability, learning and equilibrium in games with strategic complementarities, *Econometrica* 58, 1255–1277.
- Mookherjee, Dilip, and Barry Sopher, 1994, Learning behavior in an experimental matching pennies game, *Games and Economic Behavior* 7, 62–91.
- Noe, Thomas H., Michael Rebello, and Jun Wang, 2003, Corporate financing: An artificial agent-based analysis, *Journal of Finance* 58, 943–973.
- Pagano, Marco, and Ailsa Roell, 1996, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, *Journal of Finance* 51, 579–611.
- Parkes, David C., 1999, Optimal auction design for agents with hard valuation problems, *Proceedings of the IJCAI-99 Workshop on Agent-Mediated Electronic Commerce*, Stockholm.
- Pouget, S., 2007, Financial market design and bounded rationality: An experiment, *Journal of Financial Markets* 10, 287–317.
- Radner, Roy, 1979, Rational expectations equilibrium: Generic existence and the information revealed by prices, *Econometrica* 47, 655–678.
- Roth, Alvin E., and Ido Erev, 1995, Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior* 8, 164–212.
- Routledge, Bryan R., 1999, Adaptive learning in financial markets, *Review of Financial Studies* 12, 1165–1202.
- Simon, Herbert A., 1955, A behavioral model of rational choice, *Quarterly Journal of Economics* 69, 99–118.
- Walker, Mark, 1984, A simple auctioneerless mechanism with Walrasian properties, *Journal of Economic Theory* 32, 111–127.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.