Available online at www.sciencedirect.com

SCIENCE DIRECT®

Games and Economic Behavior 47 (2004) 325–352

GAMES and Economic Behavior

www.elsevier.com/locate/geb

An experiment on learning with limited information: nonconvergence, experimentation cascades, and the advantage of being slow

Eric Friedman,[a] Mikhael Shor,[b,*] Scott Shenker,[c] and Barry Sopher[d]

[a] *School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853, USA*
[b] *Owen Graduate School of Management, Vanderbilt University, 401 21st Avenue South, Nashville, TN 37203, USA*
[c] *International Computer Science Institute, Berkeley, CA 94794, USA*
[d] *Department of Economics, Rutgers University, New Brunswick, NJ 08903, USA*

Received 26 November 2001

Abstract

We present the results of an experiment on learning in a continuous-time low-information setting. For a dominance solvable version of a Cournot oligopoly with differentiated products, we find little evidence of convergence to the Nash equilibrium. In an asynchronous setting, characterized by players updating their strategies at different frequencies, play tends toward the Stackelberg outcome which favors the slower player. Convergence is significantly more robust for a "serial cost sharing" game, which satisfies a stronger solution concept of overwhelmed solvability. As the number of players grows, this improved convergence tends to diminish, seemingly driven by frequent and highly structured experimentation by players leading to a cascading effect in which experimentation by one player induces experimentation by others. These results have implications both for traditional oligopoly competition and for a wide variety of strategic situations arising on the Internet.

JEL classification: C7; C92; D83

Keywords: Learning; Limited information; Serial cost sharing; Cournot; Internet pricing; Experiments

* Corresponding author.
E-mail addresses: friedman@orie.cornell.edu (E. Friedman), mike.shor@owen.vanderbilt.edu (M. Shor), shenker@icsi.berkeley.edu (S. Shenker), sopher@econ.rutgers.edu (B. Sopher).

0899-8256/$ – see front matter
doi:10.1016/S0899-8256(03)00183-0

1. Introduction

We present the results of an experiment on learning in a continuous-time low-information setting. The setting is characterized by players having no information about the structure of the underlying game. Subjects were not informed about the structure or the stability of the payoff matrix. Participants did not know how other players' actions affect their payoffs or even that other players exist. In each period, a subject knew only the action that she selected and the resulting payoffs that she obtained. In this environment, we find that convergence is problematic even for dominance-solvable games. Interestingly, players who react more slowly may gain a strategic advantage, earning higher payoffs than players who update their actions more often. We propose that the culprit behind these non-convergence results lies in the manner in which players experiment. In contrast with many learning models in the literature, we find experimentation to be both profuse and highly structured. Further, in our low-information setting, players are unable to distinguish between sources of variation in their own payoffs—be it another player's experimentation or a change in the environment itself, perhaps caused by network delays or demand shocks. Thus, experimentation by one player may trigger experimentation by others, leading to "experimentation cascades," hindering convergence.

Although well-structured experiments with stylized notions of uncertainty have proven useful for answering many questions, real-world situations are often far more complicated and messy. For example, a person may be uncertain about how many other players are in the game, how her decisions map into payoffs (never mind the other players' payoffs), what the distribution of states is, and whether the distribution is stationary. In this paper, we are interested in the specific question: How do players learn and is equilibrium achieved in games where players know essentially nothing about the structure of the game except their own strategy space?

To model uncertain dynamic environments, our experiment proceeds in real time, provides almost no information to participants about the game's structure, and allows for asynchrony, reflecting that dynamic environments rarely have periodic, simultaneous moves. We consider three different games, each with up to six players. The first, a simple monopoly game, allows us to examine learning behavior in a non-interactive setting. The second game is a dominance-solvable Cournot oligopoly. The third is a serial cost sharing game (Moulin and Shenker, 1992), solvable by the iterated deletion of strictly overwhelmed strategies (Friedman and Shenker, 1996, 1997), a much stronger condition than dominance solvability.

Our results show that play in the Cournot game does not converge with a small number of players, despite the game being dominance solvable. In fact, players habitually select dominated actions. When play is asynchronous, strategy choice in the Cournot game tends toward the Stackelberg outcome rather than Cournot–Nash, suggesting that slower update intervals may serve as a form of commitment akin to a first-mover advantage. A "slow" player, who may update her strategy less often, effectively becomes the leader and receives significantly higher payoffs than "fast" players. This may suggest that Cournot-like mechanisms in a low-information environment reward slow or unresponsive play.[1] Hence, such mechanisms may provide players with incentive to manipulate the outcome through explicit stalling tactics or structural changes that slow their response time, leading to inefficient outcomes.

While dominance-solvability does not seem to imply convergence, play in the serial game appears to converge in synchronous and asynchronous settings. The serial game is also dominance solvable, but, in addition, it is solvable by the iterated deletion of overwhelmed strategies, which is a stronger notion. For example, overwhelmed-solvability implies that all generalized Stackelberg equilibria coincide with the Nash equilibrium. Loosely, a strategy is overwhelmed if the maximum payoff obtainable under that strategy is less than the minimum payoff obtainable under some other strategy (for details, see Friedman and Shenker, 1996). Since it is easier to deduce that a strategy is overwhelmed than that it is dominated when opponents' actions are unobservable and when players are unaware of the payoff matrix, overwhelmed-solvability may be a more robust solution concept in low-information games.

We found complimentary trends in both games as the number of players increased. Convergence in the serial game becomes less robust with more players, while convergence in the Cournot game actually improves. In Section 7, we suggest that this is due to statistical averaging, or noise canceling, and propose that such averaging may be a more relevant tool than traditional solution concepts even with a moderate number of players. Thus, in low-information environments, mechanisms that derive convergence properties from statistical averaging may be more robust than those based on iterated solution concepts.

We offer two explanations for why iterated solution concepts are inadequate notions for convergence of play in limited information environments. First, the data indicate that experimentation by subjects does not conform to commonly maintained assumptions of the theoretical learning literature. Experimentation by subjects is quite common, and more methodical than commonly assumed. A second explanation arises from the realization that an individual altering her strategy changes the payoffs of other players. From the standpoint of other players, a change in an opponent's strategy is indistinguishable from a change in the underlying payoff matrix as a cause of the payoff variation. The greater the variation a player can have on the payoffs of opponents, the less inherently stable the play of the game is, and the more likely that one player's change in strategy will invoke in other players a desire to re-initiate experimentation as if the underlying payoffs have changed, leading to "experimentation cascades." We propose that different games induce such instability at varying rates and that games must be "stable with respect to noise" for convergence to occur. In the serial game, this stability arises from the fact that a player's payoffs are independent of demands larger than her own, while in the differentiated Cournot game, stability may occur from the statistical averaging of other players' actions.

[1] For example, Greenwald et al. (2001) have shown convergence to the Stackelberg outcome in simulations of simple learning models for related games.

2. Limited-information environments

Low-information environments are found in many economic settings. In particular, the Internet gives rise to strategic interactions in which individuals are uninformed about the number of players, their payoffs, and the stability of the underlying payoff matrix (Friedman and Shenker, 1997). In models of queuing and congestion, the number of players is often unknown to the participants (Nagle, 1987; Hsiao and Lazar, 1988; Shenker, 1994; Korilis et al., 1995). Game theoretic approaches have been applied to network issues ranging from dynamic pricing (Cocchi et al., 1993; Gupta et al., 1997) to design of "shopbots," software agents who "crawl" the web on behalf of users searching for the lowest price (Greenwald and Kephart, 1999; Kephart et al., 2000). Recent interest in distributed networks (e.g., Huberman and Hogg, 1995) and the rising importance of the Internet call for analysis of learning in settings that capture basic elements of network environments.

While behavior on the Internet is a prime motivation for our experiment,[2] the relevance of this inquiry is not confined to network environments. Firms rarely know competitors' profits and often cannot observe competitors' actions. Models of collusion (e.g., Green and Porter, 1984) incorporate state uncertainty (e.g., demand shocks) so that a low market price may be due to a member of the cartel cheating or an unusually large demand shock. More generally, a player in a game with low information is in a stochastic environment in which payoff variations may be due to a number of factors including state uncertainty and variations in the decisions of other players. However, a player has very little basis to differentiate among the different causes of variations in payoffs.

Our motivation for studying learning under low information extends beyond economic settings. Low-information environments are the simplest frameworks within which to evaluate learning in repeated games. Since many models of learning only depend on personal payoffs and actions (e.g., Roth and Erev, 1995; Foster and Vohra, 1997; Sarin and Vahid, 1999; Camerer and Ho, 1999), an experimental environment in which players know only their own payoffs and actions could be used to test the implications of these models. Low-information environments allow the researcher to focus on learning behavior by limiting strategic behavior and complex strategic and psychological issues which may arise if players are more informed.[3]

Other experimental investigations concerning behavior in limited information environments include duopoly games (Fouraker and Siegel, 1963), bargaining (Roth and Murnighan, 1982), repeated matching pennies (Mookherjee and Sopher, 1994), and asset markets (e.g., Forsythe et al., 1982; Plott and Sunder, 1988). Typically, these studies find that more information improves the chance of equilibrium play especially when the information allows players to compare payoffs. Fouraker and Siegel (1963) is a notable exception, finding the opposite effect. These studies all provided experimental subjects, even in "low information" treatments, with a good deal of information about the game in which

[2] The serial cost sharing game that we study is a formalization of "fair queuing" (Demers et al., 1990), a protocol that is used by routers on much of the Internet. This game has previously been used to understand the effect of fair queuing on the quality of Internet services (Shenker, 1994; Friedman and Shenker, 1997).

[3] For example, low-information environments minimize strategic teaching (Camerer et al., 2002; Stahl, 1999) and play driven by notions of equity.

they were involved. Subjects knew the number of players in the game. Subjects were typically uncertain about one or more *parameters* of the game being played—there might be *type* uncertainty (incomplete information about the other players' payoff functions) or *state* uncertainty (incomplete information about some factor that affects all players' payoffs)—but they knew the game's general structure.

3. Experimental design

We wished to capture the following stylized facts of a low-information environment that, we believe, represent many situations described above:

(1) the nature of the game, the number of players, and the payoff matrix are unknown to the players;
(2) the underlying game is subject to structural changes through, for example, exogenous demand shocks;
(3) players change strategies at different intervals.

Limited information, in this context, means that players only observe the outcome of their own play.

Our experiment follows the lead of Chen (2002a), whose experimental design differed from most past experiments in two respects. First, players had extremely limited information, restricted to their own history of actions and payoffs; they did not know the structure of the game, the number of other players, nor could they observe those players' actions. Second, play could be asynchronous—"fast" players may change their strategies five times as often as "slow" players. Our experiment adopts these features, and confirms Chen's results about the relative performance of the serial and Cournot games in two-player settings. We extend Chen's design by considering more than two players, discretizing the strategy space into 101 actions, and raising the level of asynchrony to 30 to observe the effect of significantly larger asynchrony. However, our work differs from Chen in one important respect. Play in our setting is in (essentially) continuous time rather than in several discrete intervals. Rapid updating allowed as many as 3000 periods of play in a 50-minute experiment, which is substantially greater than previous experiments and allows a more detailed analysis of behavior.[4] Traditionally, experiments on learning span between ten and one hundred periods of play. We instituted a greater number of periods for three reasons. First, we wanted to give subjects the greatest chance to converge. Second, to test the effects of structural changes in the underlying game, enough periods for each set of parameters was desired. Third, we were interested in learning on the Internet, for which this game is a reasonable analogue.[5]

It is natural to imagine that in continuous-time games, not all players will be changing or updating strategies at the same rates. Players may choose a new action at any time, but the

[4] See Binmore et al. (2001) for a similar high-speed environment, allowing for over 1000 periods.

[5] Many adaptive algorithms on the internet, including those employed in audio and video conferencing, face precisely the environment that we are studying.

experimental design constrains such decisions to take effect at one, two, or thirty second intervals. Asynchrony is achieved by matching players of different update intervals.

We studied three different games: a monopoly game, a Cournot oligopoly, and a serial cost sharing game. A simple monopoly game allows us to analyze learning behavior in a non-interactive setting. The second game is a symmetric Cournot oligopoly with differentiated products. It was constructed to be dominance-solvable for any number of players but is not overwhelmed-solvable. Most of the learning literature suggests that play should converge to the Nash equilibrium in dominance-solvable games. The third game is based on serial cost sharing (Moulin and Shenker, 1992). It is both dominance-solvable and overwhelmed-solvable and thus robust convergence was expected.

Given its popularity, Cournot competition is a natural choice for a dominance-solvable two-player game, but to ensure dominance-solvability for an arbitrary number of players, N, we adopt a Cournot model with product differentiation. Let demand for firm i be given by

$$p_i = a - bq_i - b\left(\frac{1}{N-1}\right)\sum_{j\neq i} q_j.$$

The resulting game is dominance-solvable. Note that for $N = 2$, the above reduces to a standard two-player linear-demand Cournot model without product differentiation. Further, for this specification of product differentiation, all of these Cournot games are essentially equivalent for all N in terms of best reply dynamics, dominance, and payoffs. This arises because the game is conceptually a two-player Cournot game, in which each player's only competitor selects her quantity "by committee," i.e. by averaging the quantity choices of all $N - 1$ other firms. Naturally, this is not a property of Cournot games in general.

The last game, serial cost sharing, has been studied previously both theoretically (Moulin and Shenker, 1992) and experimentally (Chen, 2002a). A player i chooses demand q_i, and is charged a corresponding cost share, $x_i(q; C)$, of the production costs, where $C(q)$ is the cost of servicing all demands, q. The cost share, $x_i(q_i; C)$, is computed as follows. First, demands are ordered such that $q_1 \leqslant q_2 \leqslant \cdots \leqslant q_n$. Then the cost share for the player demanding the lowest quantity is given by $x_l(q; C) = C(Nq_i)/N$. Then player 1 is removed from consideration, her cost share is also subtracted from the total cost, and the remaining cost shares are computed inductively. Define $q^0 = 0$, $q^1 = Nq_1$, and $q^i = q_1 + \cdots + q_{i-1} + (N + 1 - i)q_i$ for $i > 1$. The exact formula is

$$x_i(q; C) = \sum_{k=1}^{i} \frac{C(q^k) - C(q^{k-1})}{N + 1 - k}.$$

A key property of the cost shares is that x_i does not depend on the specific value of any demand larger than q_i. The payoff function for player i is the value of her demand minus her cost share,

$$U_i(q) = v_i(q_i) - x_i(q; C).$$

We specify these functions as $v_i(q_i) = aq_i$ and $C(q) = (b/N)q^2$ for the experiment, where the parameters a and b vary by treatment. Note that this specification implies symmetry.

4. Experimental methodology

Computerized experiments were written in the Java language, and run within a Web browser.[6] Since most students have an intuitive understanding of the behavior of the World Wide Web, and understand that factors unobservable to them can affect the web's performance (such as the ability to download a file), this provided a natural experimental environment.

Each player had an update interval of either one, two, or thirty seconds. The update interval determined the length of a period for that player, that is, how often a player's change of strategy would take effect. A clock in the upper left corner of the screen notified the player when a strategy decision would next take effect by completing one full cycle each update interval. Hence, players with one second update intervals may change strategies every one second, approximating a continuous time game about as closely as possible, given current computational limitations on the communication between subjects' computers and the calculation of payoffs.

Subjects were given little information about the structure of the experiment. They were informed that the selection of a "channel," using a slider provided in the user interface, would somehow affect the payoff received for that period. They were also notified that structural changes may occur on distributed networks, such as the internet, due to server traffic, network failures, delays, etc. Subjects were not informed about the nature of the game, the underlying payoffs, the timing of changes in the games, or the number of players. Subjects were provided with only their own payoffs at the end of each period. Payoffs were calculated every second for each player. Slow players, with periods of thirty seconds, saw only the average payoff over the last thirty seconds.

In addition to the clock and channel-selection slider, the user interface (Fig. 1) provided a history of the player's payoffs on four graphs. One graph contained the entire history of payoffs. A second graph provided a "blow-up" of payoff information, showing only the most recent five periods. The player could also toggle a scatter plot between two additional graphs, indicating the average payoffs over the course of the experiment or the most recent payoffs for each channel. The slider and all graphs were color-coded by channel to make the information easier to process.

Each subject was randomly assigned to a group of between two and six subjects. All instructions, included in Appendix A, were given on the computer, and covered the use of the user interface and the aim of the experiment. Students were given as much time to review the instructions as they needed, though in only one case did reading the instructions take longer than 10 minutes. The instructions were also viewable during the game. The length of the experiment was fifty minutes, not including instructions. This allowed for 3000 periods for one-second players and 100 periods for players with update intervals of thirty seconds.

For the first ten minutes, subjects, unknown to them, were faced with a simple monopoly game with linear demand. At seven minutes, the parameters of the model changed to study the effect of structural changes on learning behavior. For the remaining forty minutes,

[6] The experiment was written in Java 1.2 and run within Sun Microsystems HotJava 3.0 Web browser. Java programs run independently of the operating system or browser used, making the experiment portable.

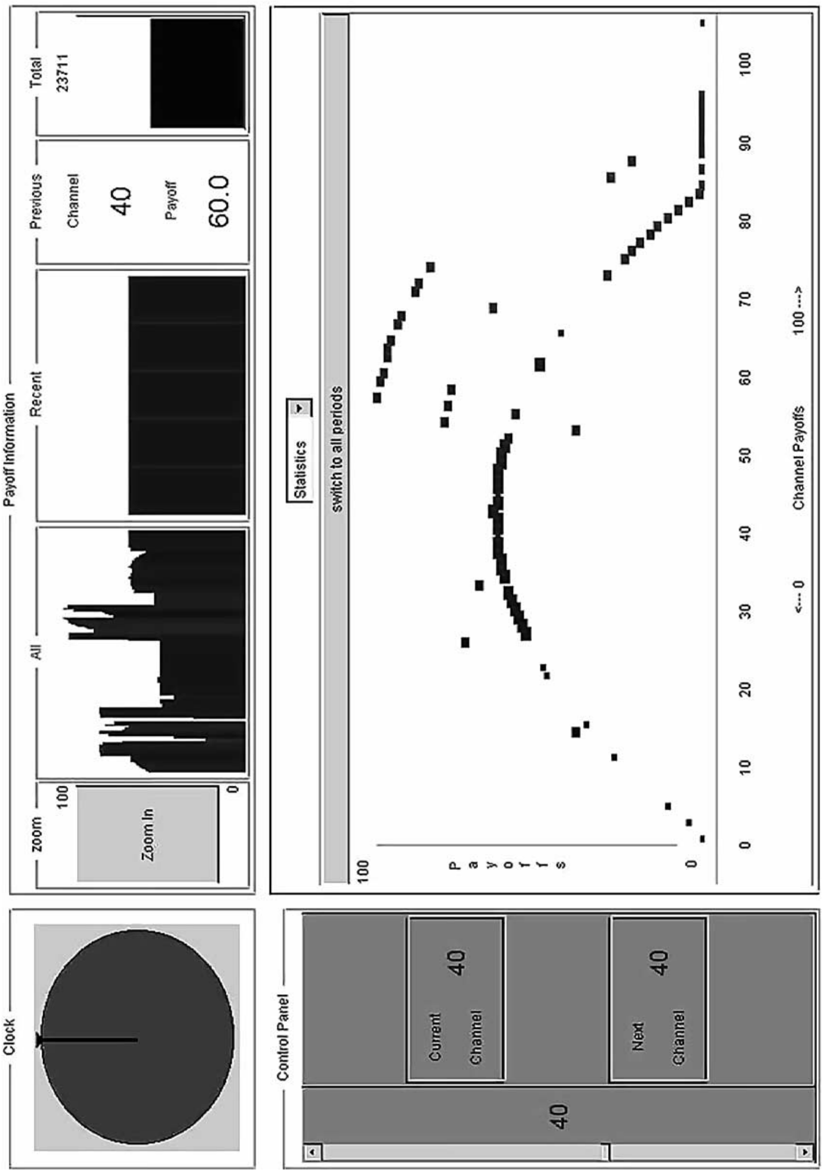

Fig. 1. The user interface.

Table 1
Summary of treatments

Time (min)	Players	Game	
0–7	1 player	Monopoly; parameter set I	
7–10	1 player	Monopoly; parameter set II	
10–40	2 or 3 players Synch. or asynch. 5 or 6 players synch.	Cournot; parameter set 1	Serial; parameter set 1
		OR	
40–50	2 or 3 players Synch. or asynch. 5 or 6 players synch.	Cournot; parameter set 2	Serial; parameter set 2

Notes. In synchronous treatments, all players had "fast" (one or two-second) update intervals. Asynchronous treatments included one "slow" player who updated at thirty-second intervals.

subjects were equally divided between the serial cost sharing game and the Cournot game. Again, the parameters of each game were changed for the last ten minutes of the experiment. Each group participated in either a synchronous or an asynchronous version of their game, yielding a total of ten treatments (Table 1). All synchronous treatments consisted of "fast" players, i.e. subjects with update intervals of either one or two seconds. Asynchronous treatments consisted of one slow player, with an update interval of thirty seconds, and the remaining players were fast players. Hence, every player participated in the monopoly section of the experiment. Then, each player participated in *either* the Cournot or the serial game for the duration of the experiment.

We began the experiment with the monopoly game for several reasons. First, the monopoly game is a special case of both the serial and Cournot games with one player. Hence, the change may be viewed not as a sudden transformation of games but simply the addition of more players. In the Cournot game, for example, the change from monopoly is akin to a new market entrant, whose quantity is determined by committee in the three and six player treatments. Second, since we conjecture that the nature of experimentation contributes to the lack of convergence, we wanted to observe experimentation in a simple environment independent of the complications of interdependent strategic response. Although it is possible that this monopoly phase affected later play, it was unlikely to have a lasting effect (Shor, 2002) and would impact the Cournot and serial games in a similar manner.

Experiments were conducted from a subject pool of undergraduate students at Rutgers University and the University of Amsterdam.[7] Experiments were run in English at both locations. Due to differences in the setup of the computer networks, player intervals were either one second or thirty seconds in the Rutgers sessions, and either two seconds or thirty seconds in the CREED sessions. Subjects were informed of their payoffs in "units" which were converted to dollars at a rate of $0.01 \cdot$ (length of period) cents per unit, and to Dutch guilders at $0.02 \cdot$ (length of period) cents per unit. Average earnings were \$24 (≈ 50 guilders), including a \$5 (10 guilder) participation fee.

The parameters for each game (monopoly, serial, Cournot) were selected to facilitate comparisons and eliminate biases created by different games yielding dissimilar payoffs at

[7] The number of subjects that participated in each treatment is in Appendix A.

equilibrium.[8] The first set of parameters for each game was chosen so that the equilibrium strategy is 40 and the equilibrium payoff is 60 for each player. The second set of parameters ensures an equilibrium strategy of 60 and payoffs of 50 (100 for the monopoly game). In each treatment, the serial and Cournot games are symmetric. Specific parameter values are given in Appendix A.

5. Results

5.1. Monopoly

As expected, seven minutes was generally ample time for players to converge to the equilibrium of a simple monopoly game in real time. Most fast players learned the equilibrium within two or three minutes. Slower players, who can update only every thirty seconds, failed to stabilize on any one strategy, and generally did not converge to or near the equilibrium in the time allotted. This suggests that fourteen periods is not enough to learn the equilibrium even in a setting as stable and as simple as the monopoly game while 100 periods is typically more than adequate.

The monopoly treatment highlights the nature of experimentation. Since the environment may change, subjects experiment occasionally to investigate the underlying payoffs. The learning literature often maintains assumptions in establishing convergence properties of various learning rules, including

(1) subjects experiment with some small probability depending only on the period, and often vanishing slowly with time; and
(2) experimentation is not correlated between players.

The first assumption is often represented by experimentation following an independent and often identical distribution in each period. Further, convergence results are generally presented as this probability vanishes. We find little support for the first assumption. The second assumption is addressed later.

Result 1. *Subjects experiment frequently and in a methodical, autocorrelated fashion.*

Even though almost all fast players learned the equilibrium within the first two minutes, experimentation persisted. Even in the last two minutes of the first monopoly treatment, we found that fast subjects deviated from the equilibrium of 40 by more than 10 (i.e., playing a strategy less than 30 or more than 50) approximately one out of every seven periods,

[8] If opponents are playing the equilibrium strategy, the Cournot and serial payoff curves coincide for strategies lower than the equilibrium, but Cournot is steeper for strategies above the equilibrium strategy. Hence, deviations above the equilibrium in the Cournot game will result in larger decreases in payoffs than the same deviations in the serial game. To the extent that a steeper payoff surface gives feedback that leads subjects to the maximum payoff more quickly, equilibrium in the Cournot game may be "easier to learn" (see the conjecture of Harrison, 1989).

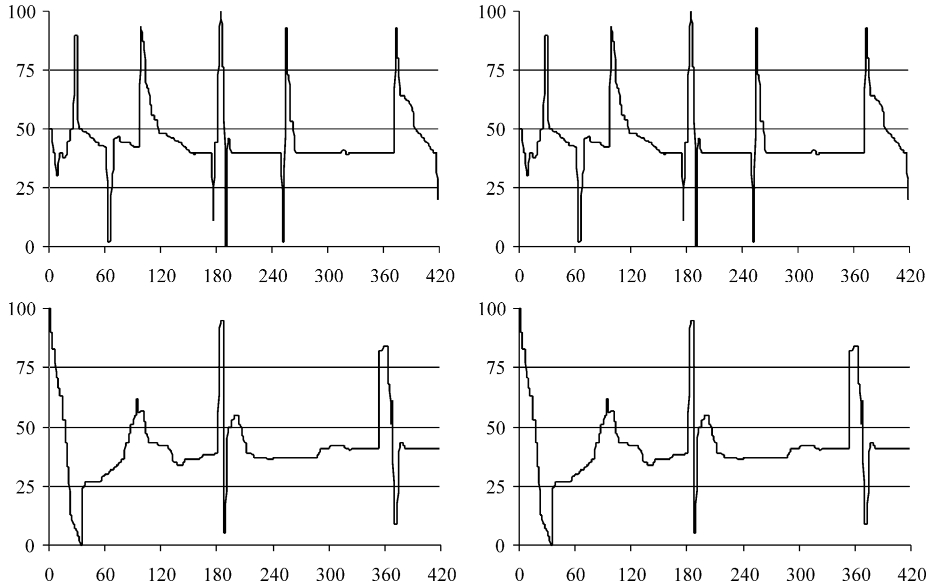

Fig. 2. Arrhythmic heartbeat patterns: path of play of selected subjects for the monopoly game exhibiting periodic and methodical experimentation.

Table 2
Experimentation by fast subjects in the monopoly game (during the last five minutes of parameter set 1)

	Distance from equilibrium for experimentation			
	1 or more	3 or more	5 or more	10 or more
Probability of experimentation in period $t+1$ following experimentation in period t	0.86	0.82	0.78	0.69
Probability of experimentation in period $t+1$ following no experimentation in period t	0.16	0.12	0.09	0.06
Percentage of subjects for whom no autocorrelation is rejected at 10%[a]	76%	83%	84%	70%

[a] If E is the set of experimentation strategies and N is the set of near-equilibrium strategies, we test for $\Pr\{s_{t+1} \in E \mid s_t \in E\} = \Pr\{s_{t+1} \in E \mid s_t \in N\}$.

reflecting substantial experimentation. Instead of exhibiting independent probabilities of experimentation in every period, subjects appeared to enter occasional "experimentation phases." Often, this involved systematically sampling to the extreme ranges of the strategy space, leading to patterns of play resembling "arrhythmic heartbeat patterns" (Fig. 2). The regular occurrence of such patterns suggests that experimentation is highly autocorrelated. For varying definitions of "experimentation" (a strategy more than 1, 3, 5, or 10 from the equilibrium), subjects are much more likely to experiment in a period following another period of experimentation than in a period directly after the subject played near the equilibrium, suggesting methodical experimentation and periodicity (Table 2).

While autocorrelated experimentation generally does not negate convergence results in learning theory in single-player environments, it does have implications for multi-player games. Since players cannot view opponents' strategies, they may confuse experimentation by opponents with a change in underlying payoffs. This may have a cascading effect, inducing further experimentation by other players, leading to complex interactions between subjects' experimentation. Even a brief experimentation period by one player can result in opponents initiating experimentation since, as the following result suggests, subjects react quickly to changes in their payoffs.

Result 2. *Subjects are responsive: they react quickly to changes in the environment, even if their current near-equilibrium payoffs are not affected*.

Many models of learning are not responsive, since experimentation rates vanish over time. While we do not know of any experiments that test whether this assumption is warranted, Friedman and Shenker (1996) emphasize the importance of responsiveness in limited-information games in nonstationary settings. Our subjects appear to be responsive, constantly experimenting and quickly adjusting to changes in the underlying game. When the monopoly game changes from parameter set 1 to parameter set 2, fast players deviate from their former strategies in an average of 18 seconds. Figure 3a presents the average strategy choice of all fast players for the monopoly portion of the experiment. The graph indicates that when the equilibrium changes, players respond quickly.

To test the robustness of players' responsiveness, we ran sessions in which, after five minutes of the monopoly game, the parameters changed, but this only affected the payoffs for strategies substantially distant from the former equilibrium (Fig. 4).[9] Hence, if players converge to a "best" strategy, eventually discarding all others, they may not be responsive to such a change. However, most players, despite converging to the equilibrium of the monopoly game, quickly discovered the change in parameters (Fig. 3b). Since payoffs at equilibrium were unaffected by the parameter change, response times were slower than when such a change was immediately felt. However, players experiment often enough (and broadly enough) to realize the change, and the figure suggests that convergence to the new equilibrium occurs within three minutes of the change.

5.2. Synchronous runs

Both the serial and Cournot games have a unique Nash equilibrium and are dominance solvable. Thus, a wide range of learning theories would predict convergence to the Nash equilibrium in synchronous play, including Bayesian learning (Kalai and Lehrer, 1993), adaptive learning (Milgrom and Roberts, 1991), fictitious play (Robinson, 1951), variants of evolutionary learning (Kandori et al., 1993; Young, 1993), reasonable learning (Friedman and Shenker, 1997), and calibrated learning (Foster and Vohra, 1997). In our environment, play in the Cournot game neither converged nor approached the equilibrium, even with only two players.

[9] The equilibrium of parameter set 1 is 40. When parameters changed after five minutes, the payoffs remained the same unless a player chose a strategy above 45.

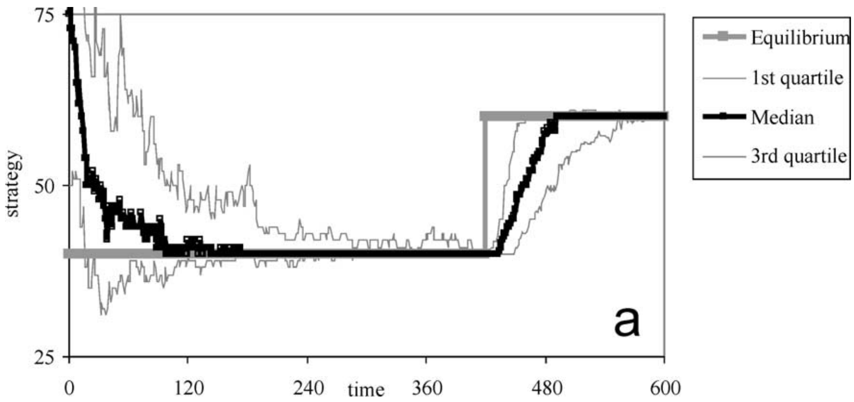

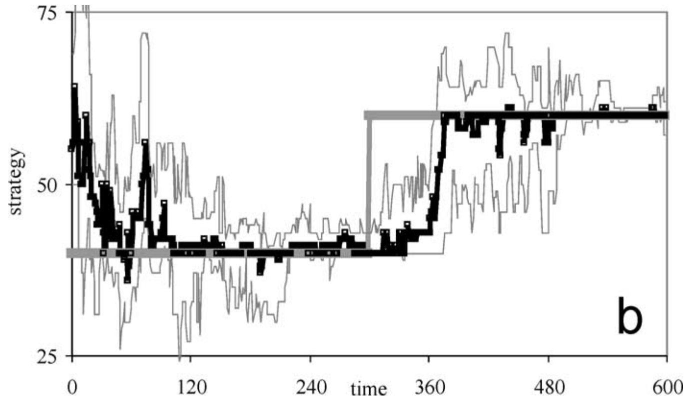

Fig. 3. Subject responsiveness to a change in the underlying algorithm for (a) Treatment I, when payoffs change at equilibrium and (b) Treatment II, when payoffs change only for strategies substantially distant from the former equilibrium.

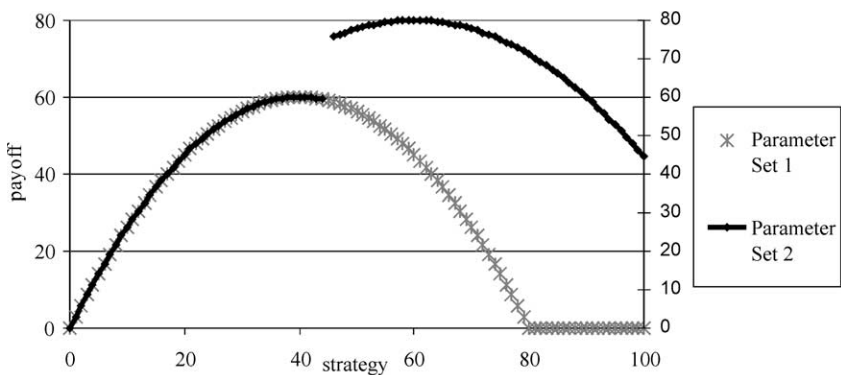

Fig. 4. Payoffs from monopoly algorithm to test responsiveness.

Result 3. *Subjects fail to converge in synchronous Cournot games with a small number of players.*

Since experimentation by subjects rarely permits strict convergence in a laboratory setting, we consider how frequently subjects chose a near-equilibrium strategy. For example, define near-equilibrium play as all strategies within 10% of the equilibrium strategy of 40, i.e. between 36 and 44, resulting in payoffs within 12% of the equilibrium payoff. The average number of periods of near-equilibrium play in the two-player Cournot game is 558 of 1800, or 33%, compared to 67% for the serial cost sharing game. Since one would not expect near-equilibrium play initially, time should be allotted for subjects to learn. We concentrate on the last ten minutes of the first parameter set, permitting twenty minutes for learning the equilibrium. Considering the frequency of near-equilibrium play for varying definitions of "near" (Table 3), the serial game appears to display strong convergence properties while the Cournot game exhibits little equilibrium play.[10]

One explanation for subjects' failure to converge in Cournot games is rooted in bounded rationality arguments. Convergence to equilibrium requires a large number of iterations of best replies or elimination of dominated strategies, yet people may be incapable of such iterations beyond a few rounds, even in settings with full information. A more dynamic, and perhaps relevant version of this point is made by Friedman and Shenker (1997) who show that if play is synchronous, then a strictly dominated action (with respect to those played with non-negligible probability by the other players) has a lower expected payoff than the strategy that dominates it. Thus, elimination of dominated strategies requires no introspection and should follow from the assumption of utility maximization. One can then show that this argument iterates and stochastic optimizing learners should iteratively remove dominated actions, if the noise is sufficiently low.

We do not find support for either of these explanations as subjects do not appear to act as if they have successfully mastered even the first step of a best reply or elimination of dominated strategies.[11] Specifically, players continually select strategies that are not a best

Table 3
Percentage of near-equilibrium plays in the two-player synchronous game (during the last ten minutes of parameter set 1)

Distance from equilibrium (in strategy space)	Game		p-value[a]
	Serial (%)	Cournot (%)	
10	93	53	0.0000
5	88	35	0.0000
1	64	15	0.0000

[a] One-tailed t-test. For near-equilibrium play, we test H_0: % serial = % Cournot.

[10] Interestingly, both two-player games are supermodular with a proper redefinition of the strategy space. Despite some intuition that strategic complementarities enhance convergence (see Chen, 2002b), this is not an adequate explanation for our data.

[11] Although for the latter explanation, it is reasonable to assume that the noise is sufficiently large that players are not able to optimize effectively. The signal to noise ratio does not seem to justify this argument unless we assume some version of bounded rationality in players' abilities to optimize in a stochastic environments.

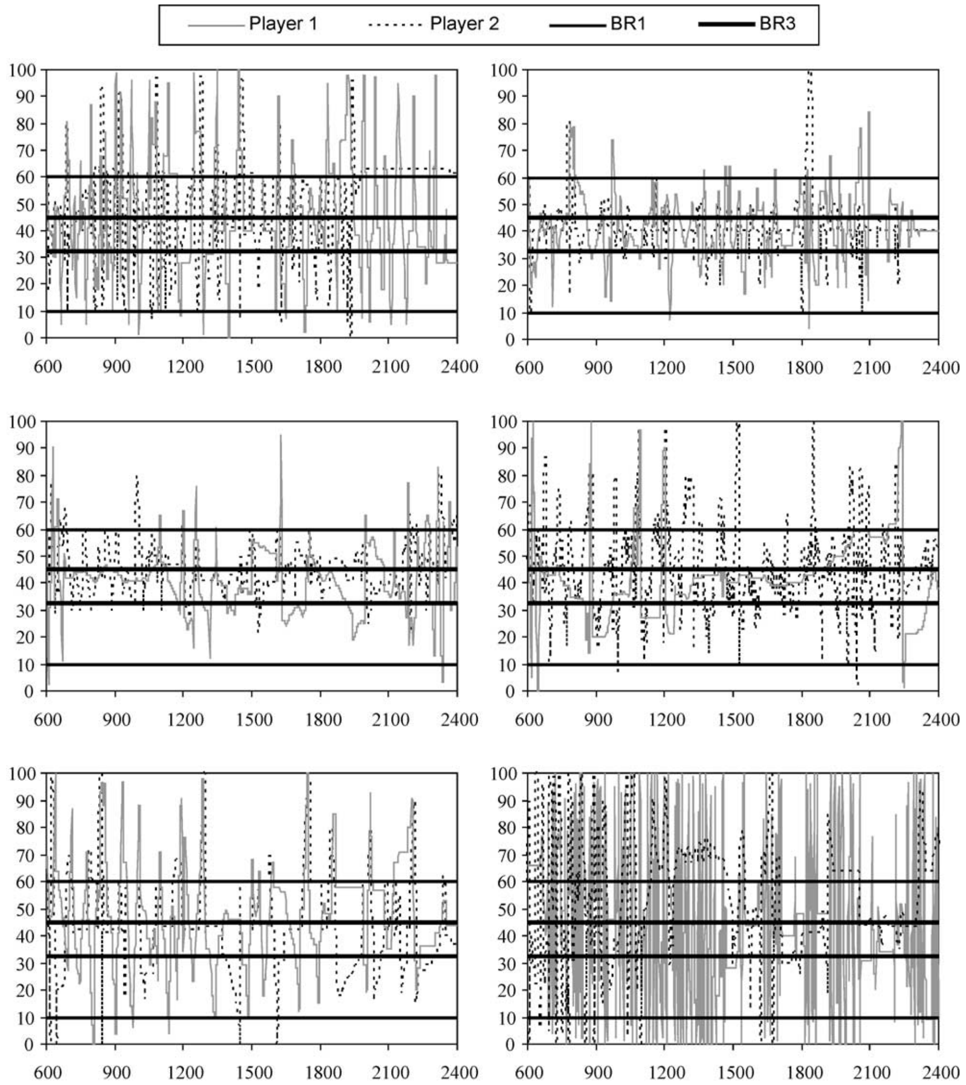

Fig. 5. Play in all two-player synchronous Cournot games. BR1 and BR3 represent the first and third iteration of the elimination of strategies that are never a best response. Time in seconds is plotted along the horizontal axis and strategy chosen is plotted on the vertical axis.

reply to any possible strategy of their opponent. In the context of the Cournot game, this implies that players do not successfully eliminate dominated strategies. We present graphs of each of the six pairs of subjects in the synchronous two player Cournot games (Fig. 5) with the first and third iterations of best reply dynamics denoted by bands. None of the pairs converges within three iterations of best replies, and only two pairs appear to remain habitually within one iteration.

The observation that players converge in the monopoly game but essentially fail to eliminate even strategies which are never a best reply in Cournot casts doubt on the use of iterated best reply or dominance concepts in analyzing learning under low information. Since from the standpoint of bounded rationality, mastering the monopoly game requires the same level of rationality as the first elimination of strategies that are never a best reply, subjects' inability to calculate or reason does not explain the lack of convergence in the Cournot game.

Result 4. *Play tends to converge to equilibrium in two player synchronous serial games*.

Most subjects played very close to the equilibrium for the last half of the thirty-minute treatment (Fig. 6). This result is similar in spirit to Chen (2002a), who finds a higher proportion of Nash equilibrium play in a serial cost sharing game than an average-cost game, similar to Cournot. Table 3 confirms that players learn the equilibrium in the two player serial games and play it consistently. However, convergence is not occurring rapidly. The graph suggests significant non-equilibrium play for the first fifteen minutes, representing 900 periods of play.

Result 5. *As more players are added, convergence in the synchronous serial game is less robust*.

The three-player serial cost sharing game also exhibits a significantly larger amount of near equilibrium play than the three-player Cournot (Table 4). However, the serial game appears to do worse in the three player game than in the two player version, while the opposite may be concluded about the Cournot game. A possible explanation, rooted in statistical averaging and the diminished effect of "experimentation cascades," is provided in Section 6. In three-player treatments, the convergence properties of the serial cost sharing game relative to Cournot remain. The Cournot game with product differentiation still fails to converge, while subjects in the serial treatment begin to play

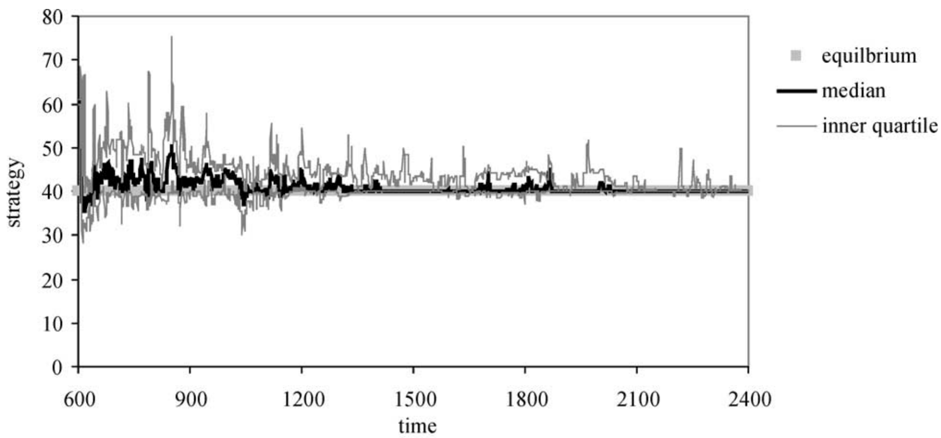

Fig. 6. Path of play in the synchronous two-player serial game.

Table 4
Percentage of near-equilibrium plays in the three-player synchronous game (during the last ten minutes of parameter set 1)

Distance from equilibrium (in strategy space)	Game: Serial (%)	Game: Cournot (%)	p-value[a]
10	84	70	0.0000
5	71	44	0.0000
1	46	15	0.0000

[a] One-tailed t-test. For near-equilibrium play, we test H_0: % serial = % Cournot.

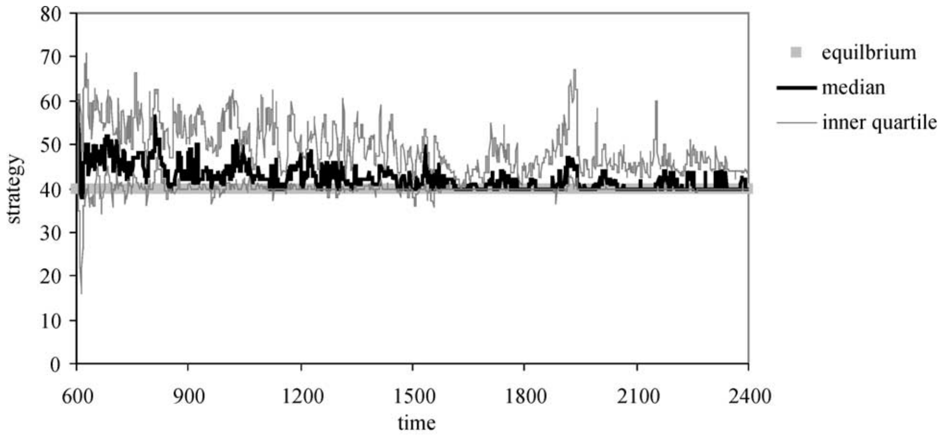

Fig. 7. Path of play in the synchronous three-player serial game.

close to the equilibrium towards the end of the thirty-minute treatment (Fig. 7). It is clear that convergence is slower and less apparent than in the two player games.

The relative advantage of the serial cost sharing game continues to vanish as more players are introduced. In Section 7, we consider five- and six-player versions of each game and offer an explanation for the improving behavior of Cournot rooted in the strategic equivalence of the games for any number of players. From the standpoint of a player, the only difference between the two and three player version, for example, is that the opponent's play in the three player version is the average of the choices of two people. This averaging implies less volatility in opponents' play as more players are introduced, and hence a more stable environment, aiding the learning process.

5.3. Asynchronous runs

When asynchrony is introduced, convergence to Nash equilibrium even in dominance-solvable games has been questioned (Friedman and Shenker, 1996, 1997). In our framework, a slow player, whose strategy changes takes effect every thirty seconds, is matched with one or two fast players who may update strategies in real time. Asynchronous play, more than simple repetition of a game in normal form, can potentially transform the game into what is essentially a repeated sequential game, which can permit Stackelberg equilibria.

Asynchronous Cournot treatments fail to converge to any outcome. However, despite this failure to converge, we can still test if play resembles a Stackelberg equilibrium rather than the Nash equilibrium. As expected, fast players significantly outperform slow players in the monopoly treatments, as shorter update intervals allow for quicker learning of the equilibrium and hence more time spent earning the maximum payoff. However, this observation does not apply to the Cournot game. In fact, the following result suggests that there is an advantage to updating slowly.

Result 6. *In asynchronous Cournot treatments, slower players earn greater payoffs than fast players and play tends towards the Stackelberg equilibrium*.

As predicted by Friedman and Shenker (1996), slow players earn higher payoffs than fast players in every run of the two-player Cournot game (Table 5), on average earning 17% higher payoffs, despite the game's symmetry. While play does not converge to the Stackelberg equilibrium, the commitment implied in slower updating led slower players to select quantities larger than the symmetric Nash equilibrium, inducing fast players, or "followers," to select smaller quantities in response. Stackelberg-like results are also found in the literature on "patient players" (Fudenberg and Levine, 1989; Watson, 1993), but in our framework, long and short term players are replaced by long and short update intervals. Hence, the structure of the game imposes a permanence on the slow player's strategy which takes the place of reputation-building.

For comparison, in the serial cost treatments, the slow player outperformed the fast player in only two runs (one is significant), and in general the difference in payoffs between fast and slow players was negligible (an average of 1.8, compared with 8.4 for Cournot). As a control, we checked the performance of the players in each treatment in the monopoly game. For both groups, slow players performed substantially worse than fast players, with subjects participating in the Cournot game performing slightly better in the monopoly treatment.

Table 5
Average payoffs in the two-player asynchronous games (parameter set 1)

Run	Cournot			Serial		
number	Slow player	Fast player	p-value[a]	Slow player	Fast player	p-value[a]
1	57.36	43.80	0.0000	55.79	57.78	0.9162
2	54.61	51.69	0.0668	53.46	56.46	0.9535
3	44.24	39.37	0.0590	49.22	56.51	0.9987
4	61.66	42.55	0.0000	47.45	46.07	0.2659
5	55.59	45.96	0.0000	57.36	49.98	0.0000
6	46.30	46.26	0.4951	50.42	57.68	0.9999
Average	53.29	44.94		52.28	54.08	
Monopoly	41.61	49.64		39.64	49.58	

Notes. Nash equilibrium payoffs are 60 in both games. The Stackelberg equilibrium payoffs in the Cournot game are 68 and 34 for the leader and follower, respectively.

[a] One-tailed t-test for H_0: slow player's payoffs equal to fast player's payoffs.

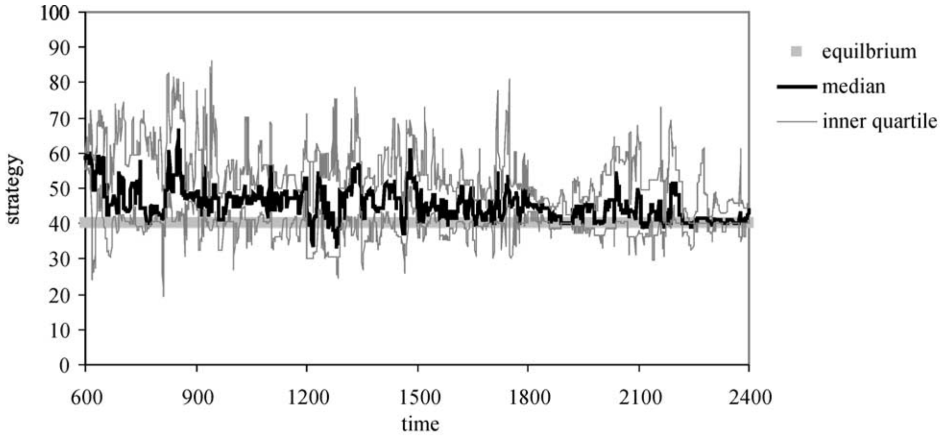

Fig. 8. Path of play in the asynchronous two-player serial game.

The serial cost sharing game continues to converge even in asynchronous settings (Fig. 8). As opposed to the Cournot game, the Stackelberg equilibrium coincides with the Nash equilibrium in the serial game. This lends support for overwhelmed solvability as a solution concept for asynchronous games, as well as synchronous games when players have limited information. However, the asynchronous treatments converge significantly slower than the synchronous treatments. Hence, in environments in which the underlying payoffs or structural elements of the game may change often, serial may fail to converge rapidly enough.

5.4. Synchronous and asynchronous comparisons

It appears that the serial game demonstrates stronger convergence in both synchronous and asynchronous treatments. We wish to summarize both cases here. Since players may continue to experiment, traditional convergence definitions stipulating that all remaining play must eventually be close to the equilibrium do not suffice. Instead, to determine the speed of convergence, we specify a definition of convergence of play following Friedman and Shenker (1997) who capture convergence as "playing close to the equilibrium most of the time." To make this notion precise, we say that a subject (ε, η)-converges at time τ if the proportion of periods after time τ which deviate from the equilibrium by more than ε is less than η. Hence, a subject (ε, η)-converges at time τ if after time τ the subject plays ε-close to the equilibrium in all but at most $(T - \tau)\eta$ periods.

Result 7. *For the two player treatments, both serial games exhibit stronger convergence than either Cournot game. Further, synchronous treatments exhibit stronger convergence than asynchronous treatments for each game.*

Figure 9 presents the percentage of individuals who (5, 0.15)-converged by time τ in each of the games. This may be viewed as the cumulative distribution function for convergence by time τ for each treatment. We adopt 15% for η, paralleling observed

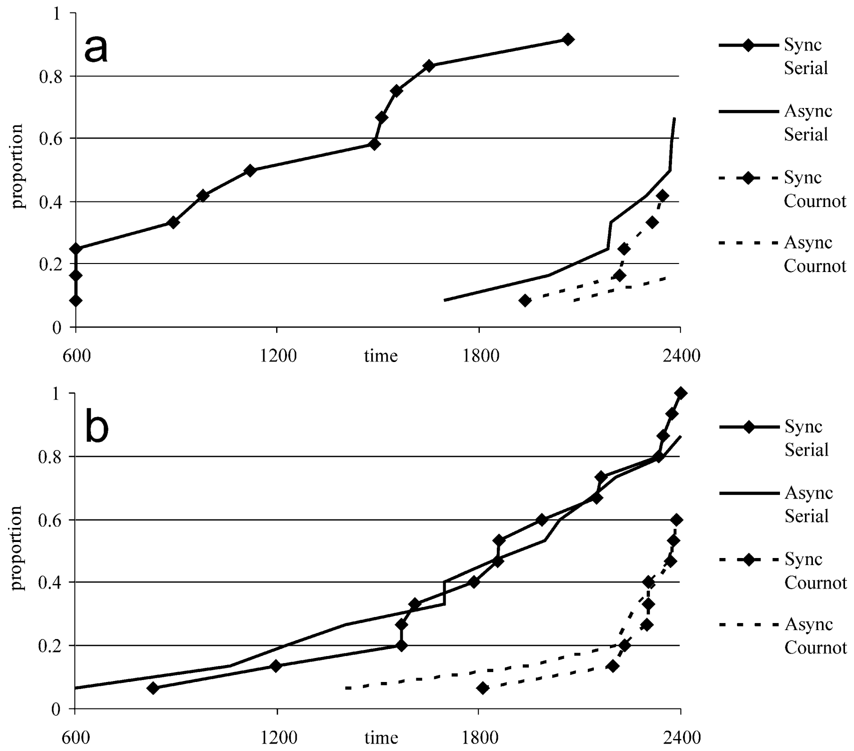

Fig. 9. Proportion of subjects who (5, 0.15)-converged in (a) two-player treatments and (b) three-player treatments.

experimentation rates in the monopoly experiment. However, the graphs are typical in the ordering of the treatments, and these results are quite robust to changes in the parameters. In the three person treatments, serial cost still performs better than Cournot, but the relatively faster convergence of the synchronous games disappears. In three-player games, the superior convergence of the serial game persists, but not as dramatically.

6. Experimentation cascades

We have suggested that some common assumptions underlying convergence results of learning models may not be warranted. Specifically, the requirement that experimentation rates go to zero does not reflect autocorrelated and frequent experimentation. This may contribute to the lack of convergence in the Cournot games, despite dominance-solvability. Intuitively, if players experiment for substantial periods of time, their opponents, observing a non-transitory change in their own payoffs, may be induced to experiment as well, introducing "experimentation cascades." Even if play is close to the equilibrium, a player's experimentation induces opponents to experiment. This leads the player to realize payoffs different than those expected which may cause additional experimentation, and so on.

To see if such correlation across players occurs, we ask if recent experimentation by one player is correlated with present experimentation by an opponent using the Granger causality test (Granger, 1969). We define experimentation, or variability of play, as the variance of play in a ten second interval. In eleven of twelve synchronous two-player runs, experimentation by at least one player Granger-causes experimentation by her opponent. Overall, 67% of all players Granger-cause experimentation by their opponents.[12] This implies that the methodical experimentation described earlier may lead to complicated cross-correlation between players.

7. Increasing number of players and noise stable equilibria

Our result that as the number of players increase convergence in the serial game weakens while convergence in the Cournot game strengthens was a surprise and does not appear to be consistent with any theories of learning with which we are familiar. To identify the effect of an increased number of players on convergence, we tested each game with five and six players. The serial game no longer exhibited convergence superior to the Cournot game. In this section, we present results for these additional sessions as well as a simple intuitive idea to explain the basic trend.

Result 8. *With a small number of players, play in synchronous serial games converges more robustly than play in Cournot games, but as the number increases, the relative advantage of the serial cost sharing game decreases.*

In order to obtain a simple measure of the efficiency of synchronous serial and Cournot games, we follow Chen (2002a) in adopting the mean absolute deviation from equilibrium. Again we only consider the last ten minutes of each treatment, allowing twenty minutes for learning. Deviation from the equilibrium across all players in a treatment (Fig. 10) reflects the superior convergence of the serial mechanism with two and three players and the trend that this relative superiority diminishes with more players. As we move to five and six players, play in the Cournot game is actually closer to equilibrium than play in the serial game. We used simple linear regression to calculate the slope of the mean absolute deviation for each run of the experiment as a function of the number of players (although a curve-fitting is presented in Fig. 10). The slope is significantly positive (p-value $=$ 0.00122, one-tailed) for the Cournot game and significantly negative (p-value $= 0.02075$) for the serial game. Here we provide a simple idea to explain this trend.

We hypothesize that in order for convergence to occur, an equilibrium must be "noise stable." By this we mean that a player's payoff function at equilibrium is not likely to be dramatically disturbed by other players' experimentation. One potential measure of this is

[12] Granger-causality test with 30 one-second lags at 5% significance. The same results were obtained at 1% and 10% significance levels.

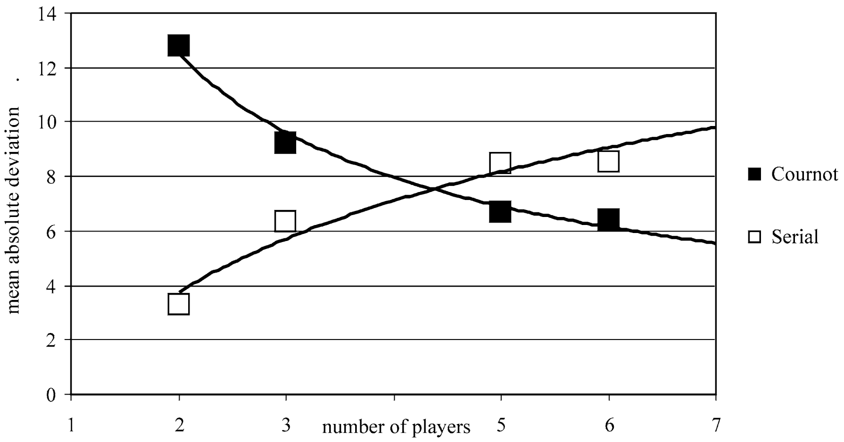

Fig. 10. Mean absolute deviation from equilibrium for all synchronous treatments.

the variance of payoffs at equilibrium due to the experimentation of others.[13] Consider the measure

$$d(q_{-i},\varepsilon)=E\big[\big\{U_i\big(q_i^*,q_{-i}(\varepsilon)\big)-U_i\big(q_i^*,q_{-i}^*(\varepsilon)\big)\big\}^2\big],$$

where q^* is the equilibrium, and $q_j(\varepsilon)$ is a random variable which with probability $(1-\varepsilon)$ is equal to q_j^* and with probability ε is a random variable X with mean q_j^* and standard deviation σ. Now we compute $d(q_{-i},\varepsilon)$ for our two games.

For the Cournot game with product differentiation, $U_i(q)=q_i(q-bq_i-b\bar{q}_{-i})$ where $\bar{q}_{-i}$ is the average of all the elements of q_{-i}. Let $\bar{q}_{-i}(\varepsilon)$ be the random variable generated by the $q_j(\varepsilon)$, $j\neq i$. Now,

$$d(q_{-i},\varepsilon)=\big(bq_i^*\big)^2E\big\{\big[\bar{q}_{-i}(\varepsilon)-q_{-i}^*\big]^2\big\}.$$

It is easy to see that

$$d(q_{-i},\varepsilon)=\big(bq_i^*\varepsilon\sigma\big)^2/(N-1)$$

and thus any measure $d(\varepsilon)$ which is some average of $d(q_{-i},\varepsilon)$ over the q_{-i}'s will be proportional to $1/(N-1)$. As N increases the equilibrium becomes less sensitive to experimentation by the players. This metric suggests that convergence in our Cournot game is particularly robust with large numbers of players. However, we do not expect this result to hold for general Cournot games with homogeneous products, which do not obey as strong of an averaging property. Thus, we view this result for the specific Cournot game only as indicative of the properties necessary for convergence and not indicating the stability of general Cournot games, even with many players.

For the serial game, such measures are not easy to compute and closed-form solutions do not exist in general due to the recursive nature of the cost-sharing rule. However,

[13] One may imagine many possible measures and we merely propose a natural and intuitive one in this context. Many other reasonable possibilities likely would yield similar results. For example, Chen and Tang (1998) measure the change in a player's best response caused by a deviation from equilibrium by an opponent.

we can solve for $d(q_{-i}, \varepsilon)$ numerically. Let X be uniform over the strategy space, and $\varepsilon = 0.15$, reflecting observed experimentation rates for the monopoly experiment. The values of $d(q_{-i}, \varepsilon)$ for the two, three, five, and six player versions of the serial game are 0.62, 0.83, 0.99, and 1.04, respectively. Note that the variance of payoffs at equilibrium is increasing with the number of players, and is concave, reflecting the observed deviation from equilibrium in Fig. 10. The ordering is robust, and is preserved for any value of ε.

Our intuition may suggest a source for less stability as more players are added in the serial game. Since the utility of a given player is independent of quantity choices larger than her own, the only relevant experimentation is that which results in an opponent selecting a strategy less than the equilibrium. Simply, as we incorporate more players, the probability of such experimentation is increasing. This effect is tempered by the fact that a single player's experimentation has less effect on other players as we add more players.

8. Conclusions

The results from these experiments have three important implications. First, convergence to the Nash equilibrium, even in a Cournot duopoly with linear demand, and despite dominance-solvability, is difficult to attain when information is very limited. In many rapidly developing industries, demand (and production) is practically unknown by the firms and opposing firms' activities are shrouded from view. In addition, market uncertainty is the norm in Internet environments as competitors update strategies in intervals measured in seconds instead of the days or months of traditional arenas. In these settings it seems likely that play can bounce around as seen in our experiments and that Nash equilibrium will be a poor predictor of behavior. In addition, there are many environments in which market participants update behavior at different rates. For example, in traditional markets, one can argue that larger firms often update more slowly than smaller ones. On the internet, update rates for real time pricing are quite varied. Thus, Stackelberg behavior might arise and, counter-intuitively, slowness might be an advantage.

Second, if standard learning models are to accurately model human behavior, they need to carefully consider the role of experimentation, which is neither infrequent nor independently distributed in uncertain dynamic settings. The frequency and autocorrelation of one player's experimentation coupled with the pursuit of another to understand her environment in light of payoff variations leads to cascades of experimentation. The highly structured nature of the experimentation observed dramatically complicates the analysis of games, but its consideration appears to be necessary for understanding economic behavior (Shor, 2002).

Third, standard (and even some nonstandard) solution concepts in game theory may need to be reevaluated for application in low-information settings, since players do not necessarily converge in the serial cost game, which satisfies almost every solution concept proposed in mechanism design. This creates new challenges for the design of mechanisms in low information environments, such as regulated oligopolies and Internet competition. We consider the development of efficient mechanisms for these settings to be an important future pursuit. In particular, such considerations are absolutely essential for establishing automated markets or markets with intelligent agents, which are becoming increasingly important on the Internet.

Our results and subsequent analysis also raise a variety of open questions and natural extensions of our experiment. These include:

(1) What would be the effect of increasing the amount of information available to the players? In addition to theoretical interest in this question, it is also of practical importance for the design of the Internet (Friedman and Shenker, 1997).
(2) In games with a second mover advantage, would asynchrony result in the faster player attaining higher payoffs? We strongly suspect that this would be the case but have not tested it.
(3) Can averaging be used to define robust mechanisms for a large number of players?
(4) What is the correct notion of noise stability?

Answering these questions would significantly enhance our understanding of many important issues in low-information environments.

Acknowledgments

The authors thank Arthur Schram, Frans van Winden, Arno Riedl, Ronald Bosman, and the rest of the members of CREED at the University of Amsterdam for their hospitality; John Ledyard, Yan Chen, participants at the Economic Science Association Conference and International Congress on Game Theory, two anonymous referees, and an associate editor for insightful comments. This material is based upon work supported by the National Science Foundation under Grant No. 9730162.

Appendix A

A.1. Parameters

Three different games were used in the experiment, and each with two sets of parameters. The serial and Cournot games were symmetric. For descriptions of the Cournot with differentiated products or serial cost sharing algorithms, see text. The monopoly algorithm is derived from linear demand, with profits given by $\Pi_j = (a - bq_j)q_j$.

The parameters used for each of the treatments are given in Table A.1. Synchronous and asynchronous runs for the same algorithm and same number of players used the same

Table A.1
Parameter values for the experimental treatments

Game	Parameter set 1[a]		Parameter set 2[b]	
	a	b	a	b
Monopoly	3	3/80	10/3	1/36[c]
Cournot	4.5	3/80	2.5	1/72
Serial	3	3/80	5/3	1/72

[a] Equilibrium strategy: 40, equilibrium payoff: 60.
[b] Equilibrium strategy: 60, equilibrium payoff: 50.
[c] Equilibrium payoff: 100.

Table A.2
Number of runs for each treatment

Game	Treatment	1 player			2 players			3 players			5 players	6 players
		Rutgers	Creed	Total	Rutgers	Creed	Total	Rutgers	Creed	Total	Total	Total
Cournot	Sync				1	5	6	1	4	5	1	1
	Async				1	5	6	2	3	5	0	0
Serial	Sync				5	1	6	1	4	5	1	1
	Async				4	2	6	0	5	5	0	0
Monopoly	Slow	7	15	22								
	Fast	27	59	86								
	Responsive	22	0	22								

parameters. The number of runs of each treatment is given in Table A.2. In each case, parameter sets 1 and 2 induce equilibrium strategies of 40 and 60, respectively. Further, in the Cournot and serial games, the payoffs to each player at equilibrium are 60 and 50, respectively.

8.2. Game instructions

A.2.1. Introduction

You are about to participate in an experiment on how people learn on the Internet. Your participation today is voluntary and will not affect your grade in the class from which you were recruited, if you were recruited from a class. A number of agencies have paid for this experiment.

You should read all of the instructions that follow carefully. If at any point you have any questions, you should raise your hand and the experiment proctor will assist you.

The instructions that follow will explain the experiment in which you are participating, as well as display and explain the parts of the screen that you will need to be familiar with. You should carefully read each page of the instructions before continuing. You will be able to view these instructions during the game.

A.2.2. Motivation

The internet, as we are all aware, has revolutionized computing. However, with the ease of communication and the global availability of data comes congestion, waiting, and in general, poor network performance.

A number of technologies and algorithms exist for potentially alleviating some of these drawbacks. For instance, in the future, if a user finds that a server from which he is downloading a file has become slow, he may switch to a faster server and continue the download without interruption.

Further, some propose that more important information, however defined, should have priority over the general user, allowing, for example, time-sensitive information to be delivered faster, at the potential expense of a recreational user.

This experiment tests some of these new ideas.

A.2.3. The clock

In the upper left-hand corner, you should see a clock. This clock counts out periods in the game. Any decisions you make will only take effect at the end of a period. A period ends every time the clock completes a cycle.

A.2.4. Clock: in motion

Now, you can see the clock spinning and counting out periods. It will spin like this during the game. A cycle is [Length of Time Interval] second[s] long, which means that every [Length of Time Interval] second[s] you may change your action.

The clock signifies when a period ends, by completing a cycle.

A.2.5. Slider

To your left, you should see the channel slider. This slider will let you select the channel that you want to use during a cycle. The slider lets you select which channel you wish to be on in a given period. You move the slider either by dragging it, clicking on the arrows, or clicking above or below the slider itself. You can practice moving the slider now. Notice that as you move it, the number under 'Next Channel' changes. The 'Next Channel' becomes the 'Current Channel' whenever the clock completes a cycle. The next instruction screen will demonstrate this.

A.2.6. Slider and spinner

Note how the slider works with the spinner. As you move the slider, the number under the 'Next Channel' changes. When the clock completes a cycle, the channel that you have selected becomes the 'Current Channel'.

A.2.7. Payoff information

The selection of a channel in each period affects your payoff for that period. At the end of each cycle, you are informed of your payoffs from the previous period or cycle. This information is presented to you in four ways. The first graph will contain many of the periods, to give you an idea of the trend of your payoffs. The second graph will only show your recent periods, or cycles. You are also given the channel which you had used in the last period, and the payoff earned last period for that channel. Lastly, you can see your total payoffs.

A.2.8. Graphs

After a number of cycles are played out, the graphs may look like this. Note that channels are color-coded both on the slider and on the graphs to make it easier to recall which channels were earning which payoffs.

A.2.9. Graphs

After even more periods of the game, the graphs may look like this. The payoffs that are earned could vary throughout the game. It is quite possible that the payoffs are so low that they are hard to read or discern on these graphs. For this reason, you are provided a 'Zoom' button to the left of the graphs. Pressing on 'Zoom in' allows you to see the lower portion of the graph. When zooming in, the text on the button changes to 'Zoom out', allowing you to return to viewing the graphs normally. Note that a scale for the graphs is provided above and below the button.

A.2.10. Statistics Panel

If you are curious how each channel has been performing on average, you may look at the 'Statistics Panel'. This panel which will appear right here when you are ready to proceed shows a plot with channel numbers along the bottom and payoffs along the vertical axis. This plot shows the average payoff over the whole game for each channel that you have used. So, if you have used some channel four times over the course of this experiment, the 'Statistics Panel' will show the average payoff that the channel has earned you over those four periods. Note that the zoom button works with the statistics panel as well.

It is possible that you may be interested in recent payoffs, rather than the average over the whole game. Above the statistics panel is a button which allows you to have only the last payoff for each channel displayed. The more recently a channel has been played, the larger the size of the 'dot' for that channel.

A.2.11. Summary

Basically, in each period, you can select a channel by using the slider to your left. At the end of each period, or cycle, you will be informed of how well that channel did last period in the graphs above and the statistics panel.

It is important to recognize that performance of the internet is not always stable. A file could be downloading at good speed, or a movie is playing over a network without any noticeable problems, and then, all of a sudden, it could slow to a crawl. This could be because a server malfunctions, a router gets overloaded, or a number of other reasons. Hence, while a channel may be performing relatively well, it should not come as a surprise if it suddenly begins to perform worse.

A.2.12. Conclusion

At this point, if you are unsure of how this experiment works, you should review the instructions again. If you have any questions now or during the experiment, raise your hand and the proctor will assist you.

The payoffs given to you in the experiment can be up to 100 units per cycle. These units will be converted into dollars at the rate of [conversion rate] cents per unit.

A.2.13. Ready

If you are ready to proceed, press the 'Done' button. You will see the actual game screen. You may move the slider to select your first channel, and make sure that you are comfortable switching between the game screen, the instructions, and the statistics panel. However, the game itself will not start until everybody has finished reading the instructions and the proctor starts the game.

References

Binmore, K., Swierzbinski, J., Proulx, C., 2001. Does minimax work? An experimental study. Econ. J. 111, 445–464.

Camerer, C., Ho, T.-H., 1999. Experience-weighted attraction learning in normal form games. Econometrica 67, 827–874.

Camerer, C., Ho, T.-H., Chong, J.-K., 2002. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. J. Econ. Theory 104, 137–188.

Chen, Y., 2002a. An experimental study of the serial and average cost pricing mechanisms. J. Public Economy, in press.

Chen, Y., 2002b. Incentive-compatible mechanisms for pure public goods: a survey of experimental literature. In: Plott, C., Smith, V. (Eds.), The Handbook of Experimental Economics Results, in press.

Chen, Y., Tang, F.-F., 1998. Learning and incentive compatible mechanisms for public goods provision: an experimental study. J. Polit. Economy 106, 633–662.

Cocchi, R., Shenker, S., Estrin, D., Zhang, L., 1993. Pricing in computer networks: motivation, formulation, and example. IEEE Trans. Networking 1, 614–627.

Demers, A., Keshav, S., Shenker, S., 1990. Analysis and simulation of a fair queueing algorithm. Internetworking 1, 3–26.

Forsythe, R., Palfrey, T.R., Plott, C.R., 1982. Asset valuation in an experimental market. Econometrica 50, 537–582.

Foster, D.P., Vohra, R.V., 1997. Calibrated learning and correlated equilibrium. Games Econ. Behav. 21, 40–55.

Fouraker, L.E., Siegel, S., 1963. Bargaining Behavior. McGraw–Hill, New York.

Friedman, E., Shenker, S., 1996. Synchronous and asynchronous learning by responsive learning automata. Mimeo. Department of Economics, Rutgers University.

Friedman, E., Shenker, S., 1997. Learning and implementation on the internet. Mimeo. Department of Economics, Rutgers University.

Fudenberg, D., Levine, D.K., 1989. Reputation and equilibrium selection in games with a patient player. Econometrica 57, 759–778.

Granger, C.W.J., 1969. Investigating causal relationships by econometric models and cross spectral methods. Econometrica 37, 424–438.

Green, E.J., Porter, R.H., 1984. Noncooperative collusion under imperfect price information. Econometrica 52, 87–100.

Greenwald, A., Friedman, E., Shenker, S., 2001. Learning in network contexts: experimental results from simulations. Games Econ. Behav. 1/2, 80–123.

Greenwald, A.R., Kephart, J.O., 1999. Shopbots and Pricebots. In: Proc. IJCAI '99, Stockholm, Sweden, August 1999.

Gupta, A., Stahl, D.O., Whinston, A.B., 1997. A stochastic equilibrium model of internet pricing. J. Econ. Dynam. Control 21, 697–772.

Harrison, G.W., 1989. Theory and misbehavior of first-price auctions. Amer. Econ. Rev. 79, 749–762.

Hsiao, M.-T.T., Lazar, A., 1988. A game theoretic approach to decentralized flow control of Markovian queueing networks. In: Courtois, P.-J., Latouche, G. (Eds.), Performance '87. North-Holland.

Huberman, B.A., Hogg, T., 1995. Distributed computation as an economic system. J. Econ. Perspect. 9, 141–152.

Kalai, E., Lehrer, E., 1993. Rational learning leads to Nash equilibrium. Econometrica 61, 1019–1045.

Kephart, J.O., Hanson, J.E., Greenwald, A.R., 2000. Dynamic pricing by software agents. Comput. Netw. 32, 731–752.

Korilis, Y.A., Lazar, A., Orda, A., 1995. The designer's perspective to noncooperative networks. In: Proc. INFOCOM '95, Boston, MA, pp. 562–570.

Kandori, M., Mailath, G., Rob, R., 1993. Learning, mutation, and long-run equilibria in games. Econometrica 61, 29–56.

Milgrom, P.R., Roberts, J., 1991. Adaptive and sophisticated learning in normal form games. Games Econ. Behav. 3, 82–100.

Mookherjee, D., Sopher, B., 1994. Learning behavior in an experimental matching pennies game. Games Econ. Behav. 7, 62–91.

Moulin, H., Shenker, S., 1992. Serial cost sharing. Econometrica 60, 1009–1037.

Nagle, J., 1987. On packet switches with infinite storage. IEEE Trans. Comm. 35, 435–438.

Plott, C.R., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. Econometrica 56, 1085–1118.

Robinson, J., 1951. An iterative method of solving a game. Ann Math. 54, 296–301.

Roth, A.E., Erev, I., 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. Games Econ. Behav. 8, 164–212.

Roth, A.E., Murnighan, J.K., 1982. The role of information in bargaining: an experimental study. Econometrica 50, 1123–1142.

Sarin, R., Vahid, F., 1999. Payoff assessments without probabilities: A simple dynamic model of choice. Games Econ. Behav. 28, 294–309.

Shenker, S., 1994. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In: SIGCOMM Symposium on Communications Architectures and Protocols, London, September, pp. 47–57.

Shor, M., 2002. Learning to respond: The use of heuristics in dynamic games. Mimeo. Owen Graduate School of Management, Vanderbilt University.

Stahl, D., 1999. A horse race among reinforcement learning models. Mimeo. Department of Economics, Univ. of Texas, Austin.

Watson, J., 1993. A 'reputation' refinement without equilibrium. Econometrica 61, 199–205.

Young, P.H., 1993. The evolution of conventions. Econometrica 61, 57–84.