

Learning to play Cournot duopoly strategies

James C. Cox^{*}, Mark Walker

Department of Economics, University of Arizona, Tucson AZ 85721-0108, USA

Received 24 March 1997; received in revised form 7 October 1997; accepted 22 October 1997

Abstract

The paper reports results from experiments designed to determine whether subjects can learn to play Cournot duopoly strategies and whether their out-of-equilibrium play is consistent with the predictions of learning models. The experiments include duopolies with constant and with decreasing marginal costs, and with theoretically stable and unstable equilibria. After a few periods, subjects do play stable interior equilibria but they do not play stable boundary equilibria nor unstable interior equilibria. Subjects' out-of-equilibrium play is inconsistent with the predictions of the learning models. © 1998 Elsevier Science B.V. All rights reserved

JEL classification: C72; C92; D43; D83

Keywords: Experiments; Learning; Duopoly

1. Introduction

The concept of the Nash equilibrium, which is central to game theory and has also become central to much of modern economic theory, depends critically upon each of the players in a game actually having correct (or nearly correct) expectations about what the other players will do. Until recently, the question of how the players will come to actually have correct expectations had generally been addressed in terms of a process of introspection: If each player's payoffs, beliefs, and rationality are 'common knowledge,' then each will be able to deduce that only Nash equilibrium strategies will be played by the others. Dissatisfaction with this introspection argument has inspired a number of recent papers that, instead, address the question of Nash equilibrium expectations in

^{*} Corresponding author.

terms of a process of learning. In one form or another these recent papers ask whether the players in a game can learn enough about one another's play, by playing the game repeatedly, to form expectations that are mutually self-fulfilling. Thus each of these papers postulates some kind of learning dynamics and then asks whether the dynamic process will converge to Nash equilibrium beliefs and/or Nash equilibrium play.

In this paper we report on a series of experiments designed to evaluate some of those theories of learning in simple Cournot duopolies. In the context of Cournot duopolies and oligopolies, the question of convergence to Nash equilibrium has a long history, both theoretical and experimental. Our experiments are designed, however, to use the simple, familiar Cournot setting to address several questions that are fundamental in understanding convergence and 'learning of equilibrium play' in any normal form game. Indeed, the simple learning models we investigate apply not just to the Cournot model, but to very broad classes of normal form games.

It is useful to organize the predictions, or implications, of theories of learning in games into three kinds of questions that can be answered only by confronting the theory with empirical evidence:

(1) Do the players (e.g. the subjects in an appropriate experiment) eventually play a Nash equilibrium? In particular, is the Nash equilibrium a good prediction of what one will actually observe in a Cournot duopoly? And, further, does the answer to this question depend upon the market demand function and the firms' cost functions?

(2) Is it possible to distinguish between two kinds of Nash equilibria – those that subjects actually do play, and those they do not play – in terms of the equilibria's (theoretical) stability properties? In other words, can the stability of equilibrium be used as a 'refinement' that serves as a useful predictive device?

(3) Do learning models provide any useful information about the path of play? Can they tell us anything, for example, about how long it takes subjects to learn to play a Nash equilibrium, or about the detailed path that their choices follow over time?

Consider, for example, the simplest Cournot model, in which demand is linear and each firm has constant marginal costs. The firms' reaction functions are as depicted in

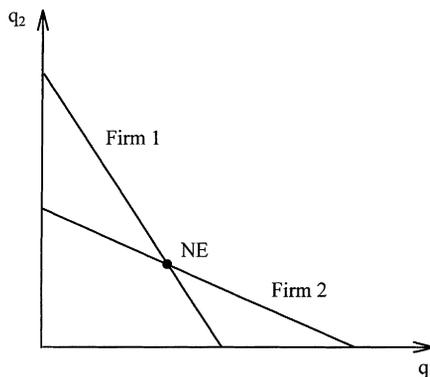


Fig. 1. Type I Duopoly.

Fig. 1: both reaction functions are linear and are also relatively flat (with respect to their own axes). Cournot’s original theory held that after a period of adjustment the firms would choose output levels that are ‘in equilibrium’ (the Nash equilibrium, in modern terminology) – namely, the pair of quantity choices that lies at the intersection of the reaction functions. Cournot further justified his equilibrium theory with a simple discrete-time dynamic theory (now called ‘best-reply’ dynamics) in which each firm always ‘myopically’ chooses the output quantity that, in hindsight, would have maximized its own profit in the preceding period.

It is easy to see, as Cournot described, that no matter what pair of quantities the firms choose to begin with, the subsequent pairs of quantities will, under his best-reply theory, follow a path that converges to the Nash equilibrium. Milgrom and Roberts (1991) have shown that a very broad class of ‘adaptive learning’ models of the players’ adjustments over time (including the Cournot best-reply model) all make the same prediction: the two firms will eventually converge to the Nash equilibrium. Many experimental studies of the Cournot duopoly model have essentially confirmed that the Nash equilibrium is a reasonably good prediction of the eventual behavior of subjects under conditions of incomplete information and constant marginal cost. (Holt, 1995 reviews the experimental research.)

But what if marginal costs are not constant? If they are increasing, the reaction functions still look qualitatively as they do in Fig. 1 and therefore adaptive learning models continue to predict convergence to the Nash equilibrium. If marginal costs are decreasing, but not decreasing too rapidly (as compared with the demand function), the reaction functions again cross as in Fig. 1, and adaptive learning models again predict convergence to the Nash equilibrium. But when one or both firms’ marginal costs are declining too rapidly, the reaction functions cross in the opposite way – firm 2’s function ‘cuts firm 1’s from above,’ as in Fig. 2. Now, in addition to the ‘interior’ Nash equilibrium where the two reaction functions cross, there are also two ‘boundary’ Nash equilibria, where the reaction functions meet again on the axes. We find it useful to refer

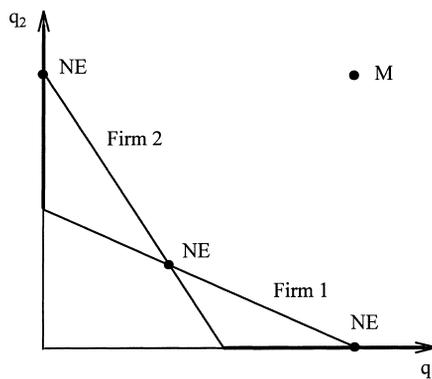


Fig. 2. Type II Duopoly

to the two kinds of duopolies depicted in Figs. 1 and 2 as Type I and Type II duopolies, respectively.

No analogue of the Milgrom and Roberts result is known for Type II duopolies – there is no general result describing the behavior of all adaptive learning models in these duopolies. However, several specific adaptive models, including Cournot's best-reply model, the model of Moreno and Walker (1991, 1994) and, if the strategy sets are discrete, the fictitious play model, all predict that play will not converge to the interior Nash equilibrium but will converge instead to one of the boundary Nash equilibria (or, for the Cournot best-reply model, that play may also alternate between output levels of zero and very high levels, without ever converging). Thus, linear Cournot duopolies constitute a theoretical example of the kind of regularity called for in questions (1) and (2) : a class of games in which we can predict, on the basis of the structure of a game and of its equilibria, whether a given Nash equilibrium can be learned (the stable one) or cannot (the unstable ones).

The experiments we report here were designed, first of all, to determine whether behavior indeed converges to the Nash equilibrium in the Type I case (even with decreasing costs) and whether it converges only to the boundary Nash equilibria in Type II duopolies, as the Cournot best-reply, fictitious play, and Moreno and Walker learning models predict (or perhaps whether it sometimes cycles, as in the Cournot best-reply model). Each model, of course, makes different specific predictions about the path of play and the rate at which play will converge.

To the best of our knowledge, no experimental studies have previously been undertaken of decreasing marginal cost duopolies, of either the Type I or the Type II variety. For this first set of experiments on the convergence properties of duopolies with decreasing marginal costs, we of course concentrate on the simplest relevant environment, duopolies in which the firms' reaction functions are linear. If the behavior we observe in this simple environment is not consistent with the theory, we will have a better chance of understanding why the theory failed than in a more complex environment. Moreover, if the theory fails in such a simple environment, there is no reason to expect that it would do better in a more complicated setting. If, however, the theory is successful, it can then be taken to more complex environments to determine the boundary between the environments in which it succeeds and those in which it fails.

In keeping with this experimental paradigm, the results we present will identify both a class of environments in which the predictions of theory are borne out, and a class of environments in which behavior is not as predicted by theory. We find that behavior does converge to interior equilibria that are stable, and that it does not converge to unstable equilibria. But when marginal costs decrease so steeply that reaction curves cross as in a Type II duopoly, where the interior equilibrium is unstable, we do not observe convergence to *any* equilibrium, not even to the stable boundary equilibria. Thus the Type I and II dichotomy is useful for distinguishing between linear duopoly environments in which behavior is or is not consistent with theory. This leaves one empirical question that must be investigated with experiments involving non-linear duopolies: is the positive result, that stable interior equilibria are played, robust to the introduction of non-linear reaction curves?

2. Linear Cournot duopolies

The general form of a linear Cournot duopoly (LCD) is as follows. The variables q_1 and q_2 denote the (non-negative) quantities the firms produce and sell; Q denotes the market quantity (i.e. $Q=q_1+q_2$); and P is the market-clearing price. The market inverse demand function is $P=\alpha-\beta Q$, where $\alpha>0$ and $\beta>0$. The firms' cost functions are $C_i = \gamma_i q_i + \delta_i q_i^2$, where $\gamma_i>0$.

The players' reaction functions are $q_1 = a_1 - b_1 q_2$ and $q_2 = a_2 - b_2 q_1$, where $a_i = (\alpha - \gamma_i)/2(\beta + \delta_i)$ and $b_i = \beta/2(\beta + \delta_i)$. (We only consider parameter values that satisfy $\alpha - \gamma_i > 0$ and $\beta + \delta_i > 0$, for $i=1, 2$; therefore each firm's reaction function is always downward-sloping.) There is an interior Nash equilibrium at $q_1^* = (a_1 - b_1 a_2)/(1 - b_1 b_2)$ and $q_2^* = (a_2 - b_2 a_1)/(1 - b_1 b_2)$.

It is easy to show that an LCD is of Type I if $b_1 b_2 < 1$ and of Type II if $b_1 b_2 > 1$. It is also easy to see (and is well known) that according to the Cournot best-reply discrete-time dynamic model the firms' choices in a Type I duopoly converge to the unique Nash equilibrium, and in a Type II duopoly they converge to one of the two boundary Nash equilibria if the firms' beginning choices lie in either the region to the upper left of the interior Nash equilibrium in Fig. 2 or the region to the lower right, and that the choices cycle between the points O and M if the beginning quantities lie in either the lower left or upper right regions.

A useful way to describe the Cournot best-reply dynamic is to say that each firm forecasts the quantity its rival will choose in the current period, and that then the firm chooses as its own quantity the output level that will maximize its own profit if its rival's quantity turns out to be as forecasted. And, of course, the quantity each firm forecasts for its rival in the Cournot best-reply model is the quantity the rival was observed to have chosen in the preceding period.

Moreno and Walker (1994) describe an alternative dynamic model in which each firm behaves as in the Cournot best-reply model except that the firm's forecast of its rival's choice is the average (the sample mean) of all the quantities the rival has chosen in the past. It is not hard to show (see, e.g. Moreno and Walker, 1991) that if the two firms behave in this way, then in Type I duopolies their quantity choices will converge to the (unique) Nash equilibrium, and in Type II duopolies the firms' choices will converge to one of the two boundary Nash equilibria. In other words, while the paths will differ in the two models, both the Cournot best-reply model and the Moreno and Walker model make the same predictions about which Nash equilibrium will be 'learned' by the firms in each type of duopoly.

Other learning models of an adaptive character have the same convergence properties as the Moreno and Walker model. We will describe one such model, the model of fictitious play, in Section 3.

3. Experimental LCDs and discrete strategy sets

The experiments were designed around three different types of LCDs: Type I LCDs in which each firm has constant marginal cost; Type I LCDs with decreasing marginal cost

for each firm; and Type II LCDs with decreasing marginal cost for each firm. For experimental implementation, each duopoly has discrete strategy sets: each firm can choose any non-negative *integer* output quantity.¹

The move from continuous to discrete strategy spaces changes the models. The Moreno and Walker model, for example, is defined for continuous strategy spaces; with discrete strategy spaces, a firm cannot generally best-respond to a forecast consisting of the sample mean of its observations because the sample mean will not typically be an integer. We redefine the Moreno and Walker model's forecasts to consist of the sample mean, *rounded* to the nearest integer. The Moreno and Walker results on convergence do not apply to this discrete model. We address this issue via the numerical simulation described in Section 5.

On the other side of the ledger, the move to discrete strategy spaces makes the theoretical results on fictitious play available to us. In the fictitious play model, each firm's forecast of its rival's output is in the form of a probability distribution, in effect a mixed strategy. (Indeed, the model was first introduced by Brown (1951) as a procedure for computing mixed strategy equilibria, not as a learning model.) The firm forecasts that its rival will choose its quantities with probabilities corresponding to the frequencies with which the rival has been observed to choose in the past.

For Type I duopolies, a theorem of Milgrom and Roberts ensures that, if firms behave according to the fictitious play model, then their choices will always converge to the Nash equilibrium. There are nevertheless two considerations that lead us to undertake numerical simulation for the fictitious play model, just as with the Moreno and Walker model. First, for purposes of experimental design – the choice of parameter values to use in creating the experimental duopolies and the number of decision periods to allow – it is important that we know *how long* the model that predicts convergence will take. Second, there are no general theoretical results about the convergence behavior of the fictitious play model in Type II duopolies. Finally, we need numerical solutions for all the models in order to be able to test their convergence *path* predictions.

4. Random matching of subjects

The central question we are investigating is whether participants in a strategic situation might learn to play one of a game's Nash equilibria by playing the game repeatedly. The learning models we are considering all assume that in a repetitive situation of this kind the players will behave purely 'reactively,' or myopically. The models assume that players will not employ strategies designed to influence the future play of the other players, such as strategies that lead to reputation building or to punishment of recalcitrant opponents, and that a player also will not take account of the fact that

¹ The discreteness introduces a conflict between uniqueness of the interior Nash equilibrium, on one hand, and, on the other, very small opportunity costs of choosing quantities near the Nash equilibrium levels. See the discussion in Section 6.

he can, by his own actions, affect the information he (or his opponents) will have available at future plays.

In designing experiments based on these ‘reactive’ theories, one must devise strategic situations in which subjects will indeed be likely to behave reactively. On the other hand, the games we are studying are duopolies, that is, two-player games, a domain in which it is natural to expect a more ‘strategic’ kind of behavior. In order to make such repeated-game behavior unlikely, our experimental design randomly matches subjects into pairs, with a new configuration of matches generated in every period. The matching protocols are structured in such a way that each subject is matched only rarely against any other particular subject. A subject does not know at which periods he has previously played against (or will in the future play against) her current opponent. As explained in Section 6, different random matching protocols were used in the two sets of experiments.

It might be argued that the attempt to induce reactive behavior is misdirected, that it is unrealistic to expect players in a ‘real’ duopoly to eschew repeated-game behavior. But we study duopolies not because we are interested only in duopolies, but because we are also interested in games with enough players that reactive behavior is plausible, or even likely. We concentrate here on two-player games only for theoretical and experimental tractability, just as we study market behavior in Edgeworth-box economies in order to understand market behavior when there are actually many participants.

5. Simulations

Since our duopolies are discrete and the subjects are randomly matched in each period, existing results for the dynamic models we wish to analyze do not directly apply. We have not attempted to work out general theoretical properties of all discrete duopolies with randomly matched players. Instead, for the specific duopolies in our experimental design, we have performed numerical simulations of the dynamic models we want to investigate. As we expected, the discrete duopolies with random matching performed qualitatively the same in simulations as theory tells us that duopolies will perform when player matchings are held constant across time, assuming players play according to the Cournot best-reply, fictitious play, or Moreno and Walker model. We provide a brief description of these simulation results.

In both constant and decreasing marginal cost Type I duopolies, play always converges, in both the Moreno and Walker and fictitious play models, to the set of choices consistent with Nash equilibrium. Unless the players’ initial expectations are quite skewed, the convergence is rapid: most players never change their choices after about the first 10 periods, and none after about the first 15 periods.

In decreasing marginal cost Type II duopolies, play always converges, in both the Moreno and Walker and fictitious play models, to one of the two boundary Nash equilibria, which one depending upon play at the first period. Most players never change their choices after about the first 10 periods, but it sometimes requires up to 50 or 60 periods before every one of the 20 players has reached her Nash equilibrium choice.

6. Experimental design

Eight experiments were conducted, each with 20 subjects participating in a sequence of several market periods, each subject playing the role of a firm in a Cournot duopoly. In each experiment, the 20 subjects were divided into two groups, consisting of 10 Type A subjects, all with the same cost function, and 10 Type B subjects, all with the same cost function, but different from the Type A cost function. In each market period there were 10 duopoly markets, each market containing one of the Type A subjects and one of the Type B subjects. In all but one of the experiments there were 30 market periods and the subjects were randomly rematched into new pairings in each period. In the remaining experiment, subjects were randomly matched in the first market period and those pairs remained fixed during that experiment's 37 periods. The subjects were not told, in any of the experiments, how many market periods there would be. Each subject was informed of his own cost function, which remained the same throughout the experiment, and of the fact that all subjects of his type had the same cost function. Subjects were not informed of the cost function of the other type of subject; they were told simply that the cost function of the other type could either be the same or different from their own type's cost function. In all the experiments, the Type A and Type B cost functions did, in fact, differ from one another. Information about a subject's own cost function was presented on screen in the form of equations and graphs of marginal costs, with an accompanying written explanation.

The demand side of the market was implemented by having the computer buy all units supplied by the duopolists, at a price determined by a market-inverse demand function which was common information to the subjects. This inverse demand function was presented to the subjects via an equation and a graph and described in the accompanying text. Individual subjects could privately use a 'profit calculator,' built into the network software, to try out conjectural quantity choices. When a subject enters into the calculator a hypothetical quantity for himself and one for the other seller in the market, the calculator reports the price that would result from those quantities according to the demand function, and also reports what the subject's own profit would be (but not, of course, the rival's profit).

The design characteristics of the eight experiments are summarized in Table 1. Note in Table 1 that the experiments can be classified into three categories, according to the

Table 1
Experiments

Experiment	Duopoly Type	Marginal Cost	Matching	Interior Equilibria
1. SC(1)	I: Stable	Constant	Random	3
2. SC(2)	I: Stable	Constant	Random	1
3. SD(1)	I: Stable	Decreasing	Random	3
4. SD(2)	I: Stable	Decreasing	Random	1
5. UD(1)	II: Unstable	Decreasing	Random	2
6. UD(2)	II: Unstable	Decreasing	Random	1
7. Bdy	II: Unstable	Decreasing	Random	2
8. UDF	II: Unstable	Decreasing	Fixed	2

duopoly type (Type I (stable) or Type II (unstable)) and the nature of marginal costs (constant or decreasing). This yields only three categories (SC, SD, and UD), because Type II duopolies cannot have constant costs.

An important objective in the design of the experiments was that, as nearly as possible, the interior equilibria be the same in each of the three duopoly markets, SC, SD, and UD, so that any observed differences in play could not be attributed simply to differences in the equilibria. This objective, however, runs up against an unavoidable difficulty caused by the combination of discrete strategy spaces, quadratic payoff (profit) functions, and the three kinds of duopoly: if the interior equilibrium is unique for each duopoly, then a subject who deviates from equilibrium play will suffer only an extremely small reduction in his payoff, compared to the absolute level of his payoff. The only alternative to this is a design with multiple, adjacent interior equilibria. We therefore conducted two series of experiments, each series including an SC, an SD, and a UD duopoly market. In Series 1 each duopoly has multiple adjacent interior equilibria: $(q_a, q_b) = (4, 9)$, $(5, 8)$, and $(6, 7)$ in SC and SD, with rationalizable strategy sets $\{4, 5, 6\}$ for Firm 1 and $\{7, 8, 9\}$ for Firm 2; and $(q_a, q_b) = (6, 8)$ and $(5, 9)$ in UD. In Series 2 each duopoly has the unique interior equilibrium $(q_a, q_b) = (3, 5)$, but the quantities $q_a = 4$ and $q_b = 6$ yield only slightly smaller payoffs (against one another, or against the equilibrium quantities), and we therefore refer to these as ‘flat-maximum’ choices.

Thus, the first six experiments listed in Table 1 are distinguished as Series 1 or Series 2 experiments. The experiments are labeled $SC(k)$, $SD(k)$, and $UD(k)$, where k denotes the series; S denotes a stable (Type I) duopoly, and U an unstable (Type II) duopoly; and C indicates that marginal costs were constant, D indicates that they were decreasing.

The remaining two experiments, labeled Bdy and UDF, applied two alternative treatments to one of the Series 1 experiments, $UD(1)$. The experiment UDF was identical to $UD(1)$ except that pairs were kept fixed throughout, instead of being randomly rematched in each market period. This experiment was designed to serve as a benchmark for the UD experiment: while there is an existing body of experimental research using Type I duopolies with rivals fixed across time – experiments against which one can compare our random-matching Type I experiments – there has been no experimental research with Type II duopolies. Experiment UDF was designed to fill this void and to serve as a basis for comparison with the results in the random-matching experiment, $UD(1)$.

The experiment labeled Bdy was motivated by a systematic feature of the $UD(1)$ data: while the subjects’ choices in experiment $UD(1)$ appear to have been ‘repelled’ by the unstable interior equilibria, they do not appear to have been ‘attracted’ by the stable boundary equilibria at which one of the subjects had to choose a quantity of 0. In order to determine whether this was simply due to subjects’ resistance to choosing the quantity zero, the Bdy experiment introduced a ‘boundary’ treatment, which moved the extreme, stable equilibria away from the $q_i = 0$ boundary of the strategy space.

The demand and cost parameters for each of the eight experiments are given in Table 2. All prices and costs were quoted to the subjects in experimental dollars (E\$), which were converted into U.S. dollars for the subjects’ earnings. Each subject was informed in the instructions of the exchange rate common to all subjects of his type, but not the rate for subjects of the other type. In $SC(1)$ and all the Series 2 experiments, exchange rates were set so that the U.S. dollar earnings at the interior Nash equilibria

Table 2
Design parameters

Series 1			Series 2		
Demand: $P=60-2Q$			Demand: $P=450-30Q$		
Marginal costs			Marginal costs		
	Type A	Type B		Type A	Type B
SC(1)	$MC_a=24$	$MC_b=18$	SC(2)	$MC_a=100$	$MC_b=40$
SD(1)	$MC_a=\max$ $(0, 29-q_a)$	$MC_b=\max$ $(0, 26-3q_b)$	SD(2)	$MC_a=\max$ $(0, 170-5q_a)$	$MC_b=\max$ $(0, 130-5q_b)$
UD(1)	$MC_a=\max$ $(0, 39-3q_a)$	$MC_b=\max$ $(0, 42-3q_b)$	UD(2)	$MC_a=\max$ $(0, 275-22.5q_a)$	$MC_b=\max$ $(0, 305-22.5q_b)$

Parameters for experiments Bdy and UDF were the same as for UD(1).

All prices and costs are expressed in Experimental Dollars. U.S. dollar exchange rates are given in Table 3.

Table 3
Exchange rates and earnings

Exchange rates ^a Experiment	Type A	Type B	Earnings ^b		
			Low	Mean	High
1. SC(1)	100	250	\$12.18	\$14.69	\$18.10
2. SC(2)	396	1020	\$31.25	\$36.64	\$42.50
3. SD(1)	100	100	\$9.28	\$14.41	\$21.59
4. SD(2)	216	660	\$17.00	\$27.61	\$40.75
5. UD(1)	100	100	\$0.02	\$9.59	\$23.59
6. UD(2)	90	240	(\$1.10)	\$22.21	\$58.44
7. Bdy	100	100	\$13.25	\$21.46	\$30.50
8. UDF	100	100	\$0.73	\$13.16	\$34.81

^aExperimental dollars per U.S. dollar.

^bDoes not include \$5 show-up fee paid to each subject.

were the same for the two subject types. In the other experiments the exchange rates were E\$100 per U.S. dollar. Table 3 gives the exchange rates in each experiment and the U.S. dollar amounts that the subjects earned. The experiments lasted $1\frac{1}{4}$ to $1\frac{1}{2}$ hours, including time for subject sign-in, instructions, and payment of earnings.

In order to minimize the chance that subjects would become bankrupt during the course of an experiment, two additional features were included in the design of the experiments and in the accompanying instructions to the subjects: (a) each subject was given an initial capital balance of E\$500, and (b) no subject could lose more than E\$30 in any period (i.e. a subject's net profit in each period was $\max\{-30, \text{Revenue}-\text{Cost}\}$). With the exception of the experiment UD(2), no subject ever came close to bankruptcy, although losses at the maximum E\$30 were not uncommon.

The procedures by which subjects were randomly matched were as follows. In the Series 2 experiments, the matching probabilities in each period t were independent of history prior to t : every configuration of A–B matches was equally likely at every period. The matching in the Series 1 experiments was carried out so as to ensure that any given Type A subject would play against each Type B subject exactly once in the first 10 periods, once in the second 10 periods, and once in the third 10 periods, in a random order

within each of the 10-period blocks, and with an additional provision that prevented any subject from playing the same other subject in two successive periods (i.e. in periods 10 and 11 or periods 20 and 21). The random matching protocol was explained to the subjects in non-technical language in the experiment’s instructions. Subjects were never given any further information from which they could determine who they had been matched with in any previous period, or who they were matched with in the current period.

7. Experimental results

In this section we describe the results from the eight experiments. We first ask whether the subjects’ behavior converged to Nash equilibrium quantity choices, and then we ask whether the three alternative learning models can successfully account for the period-by-period play of the subjects. We will be especially interested in whether observed play was significantly different in the four Type I experiments, where learning theories predict that play will be at or near the interior equilibria, from play that was observed in the four Type II experiments, where theory predicts that the interior equilibria will not be played.

7.1. Are Nash equilibrium quantities eventually chosen?

The question whether subjects eventually ‘learn’ to play Nash equilibrium is addressed in Fig. 3 and Table 4. Both the figure and the table reveal a striking difference between play in Type I and Type II duopolies. In Fig. 3, each of the line graphs corresponds to one

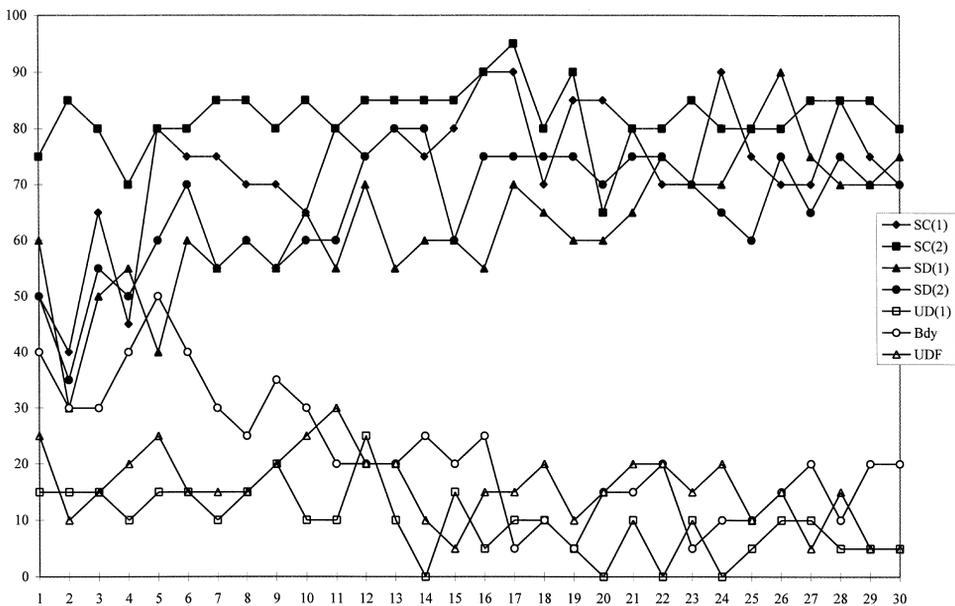


Fig. 3. Play of interior equilibrium or flat-maximum quantities.

Table 4
Experimental results: Convergence to equilibrium

Experiment	Duopoly type	Returns to scale	Interior NE or flat ^a
1. SC(1)	I: Stable	Constant	70–90%
2. SC(2)	I: Stable	Constant	80–85%
3. SD(1)	I: Stable	Decreasing	65–90%
4. SD(2)	I: Stable	Decreasing	60–90%
5. UD(1)	II: Unstable	Decreasing	0–10%
6. UD(2)	II: Unstable	Decreasing	n/a
7. Bdy	II: Unstable	Decreasing	5–20%
8. UDF	II: Unstable	Decreasing	5–15%

^a Percentage of subjects who played their part of a Nash equilibrium or flat-maximum profile in the last 10 periods.

of the experiments, and the vertical axis (the height of each of the line graphs) measures the percentage of the 20 subjects in that experiment who played their part of an interior Nash equilibrium or (in the Series 2 experiments) a flat-maximum quantity. The four line graphs describing play in the four Type I experiments (the SC and SD experiments) clearly rise over time, and after the first few periods are all well above the 50 percent level. The three graphs describing play in the Type II experiments, on the other hand, tend to be declining, and after the first few periods they are all far below the 50 percent level.

The rightmost column of Table 4 gives a summary of play during the last 10 periods of each experiment. The low and high percentages that appear in this column for each experiment are the minimum and maximum values of its line graph in Fig. 3 over the last 10 periods of the experiment. Just as in Fig. 3, it is clear in Table 4 that after subjects had accumulated some experience in the early periods, play in the Type I experiments was then predominantly at equilibrium or flat-maximum quantities (all these percentages are 60% or higher), but play in the Type II experiments was rarely at equilibrium or flat-maximum quantities (all percentages are 20% or lower).²

No equilibrium play is reported in Fig. 3 or Table 4 for experiment UD(2) because the experiment was terminated after nine periods: three of the Type B subjects had gone bankrupt by that time, and all of the other Type B subjects had lost money. Play thus seems to have clearly been converging in this experiment to one of the boundary equilibria, in which the duopoly degenerates into a natural monopoly.

The results of the experiments, as reported in Fig. 3 and Table 4, strongly suggest that the theoretical stability properties of a Nash equilibrium can serve as an effective ‘refinement,’ distinguishing equilibria that subjects can be expected to play from those that they generally will not play.

7.2. Results from the fixed-matching experiment UDF

The design of experiment UDF differed from that of UD(1) in only one respect: in UDF the subjects were matched at the beginning of the experiment and the matches then

² Note that there were 37 periods in experiment UDF, so that the figures in Table 4 refer to periods 28–37, not to periods 21–30 as in the other experiments.

remained fixed, while in UD(1) the subjects were rematched in each period. As reported in the last line of Table 4, the 5–15 percent of quantity choices in UDF that are consistent with an interior Nash equilibrium are similar to the rate of interior equilibrium play in the other Type II experiments, in which subjects were rematched in each period, and this rate of equilibrium play is far below the rate in the Type I experiments reported in the top half of Table 4.

In spite of the surface similarity in behavior, comparison of the data from experiment UDF with the data from experiments UD(1) and Bdy indicates that rematching of subjects is an important design feature in duopoly experiments: there was a significant incidence of ‘cooperative’ behavior in UDF that did not appear in the other two experiments, nor in any of the Type I experiments. The most profitable quantity choices for a pair of cooperating duopolists in this decreasing cost duopoly are for one firm to sell 15 units and the other to sell 0 units. No side payments were possible during the experiment (or after, since no subject knew the identity of her rival subject), and therefore the only way for the subjects to cooperate was to ‘take turns’ operating as monopolists, with only one subject selling a positive quantity in each period, and for the active seller in a period to choose a sales level of 15 units. Achieving this cooperation – that is, settling on a division of the spoils (with what frequency does each subject get to be the monopolist?) and also attaining the requisite common understanding of who will be active in which periods – would seem to be a formidable task: the only form of communication possible was through the subjects’ quantity choices.

A remarkable feature of subjects’ choices in the UDF experiment is evident in the detailed, disaggregated data, which we provide in Appendix A. By the end of the experiment, four of the 10 subject pairs were coordinating their play, alternating between (15, 0) and (0, 15), and a fifth pair seemed to be on the verge of achieving such coordination. It is interesting to observe *how* some of the subject pairs managed to cooperate, given the limited possibility for communication. In Appendix A we describe several patterns of successful and unsuccessful quantity signalling that can be identified in the data.

7.3. Do learning models accurately predict the path of play?

The Cournot best-reply model, the Moreno and Walker model, and the model of fictitious play are all models of ‘learning to play’ Nash equilibrium: they yield predictions of how duopolists’ quantity choices will be determined by the history of play. Thus, each model yields a prediction for each subject for each period in each experiment (after period 1). Table 5 reports full-distribution tests of consistency of the subjects’ choices with the path predictions of the three learning models. The table reports results from Kolmogorov–Smirnov (K–S) and Epps–Singleton (E–S) tests for the Moreno and Walker, Cournot best-reply, and fictitious play models.³ At the one percent significance level, data from all seven random-pairing experiments are inconsistent with the path predictions of all of the models.

³ See Conover (1980) and Epps and Singleton (1986), respectively, for explanations of these tests.

Table 5
Full distribution tests of the learning models

Experiment	Model	K–S test	E–S test
SC(1)	FP	10.84	159
	MW	7.89	141
	BR	8.43	176
SC(2)	FP	12.22	271
	MW	9.42	73
	BR	12.34	321
SD(1)	FP	8.30	126
	MW	8.06	125
	BR	8.06	160
SD(2)	FP	11.36	194
	MW	9.16	287
	BR	10.95	259
UD(1)	FP	11.17	350
	MW	14.66	474
	BR	14.91	518
UD(2)	FP	6.32	381
	MW	6.72	307
	BR	6.32	311
Bdy	FP	10.63	303
	MW	9.67	223
	BR	10.34	463

8. Conclusion

The clearest conclusion to be drawn from these experiments is the difference between the predictive value of Nash equilibrium in Type I and Type II duopolies. Fig. 3Table 4 reveal this difference in sharp relief. In the four experiments with Type I duopolies (the SC and SD experiments), in which all three learning models predict convergence to the interior Nash equilibria, most subjects chose Nash equilibrium quantities after only a few periods, and deviations from Nash equilibrium play were small. The presence of decreasing costs in SD(1) and SD(2) had little effect. In contrast, in the four Type II experiments, in which the learning models predict convergence only to the boundary Nash equilibria, subjects rarely chose quantities that were part of an interior equilibrium, and variability in choices across subjects was substantially higher than in the Type I experiments. Moreover, there is no evidence of movement over time toward interior equilibrium play in any of the four Type II experiments.

Thus, these experiments strongly suggest that the theoretical stability properties of a Nash equilibrium can serve as an effective ‘refinement’: it appears to be reasonable to expect players to eventually play stable equilibria but not unstable ones. While all but one of our experiments rematched subjects with one another in every period, the sharp distinction between the predictive value of Nash equilibrium in Type I and Type II duopolies might be expected to also hold in games with an unchanging

set of players if the number of players is large enough to induce reactive, myopic behavior, as we might expect in a market with several firms. This is clearly a question for further research.

In the experiment UDF, in which the subject pairings remained fixed throughout the experiment's 37 periods, an unexpected phenomenon appeared which seems to be rich with suggestions for both theoretical and experimental investigation. One of the subjects in almost every pair recognized that he and his rival would both fare better if they coordinated their play, one subject in each period selling zero units and the other selling the monopoly quantity of 15 units, and alternating which subject would be the monopolist. The experimental design provided the subjects with no means for communicating with one another to effect this coordination, except via their quantity choices (each of which was revealed after each play). In nearly every pair, one of the subjects appears to have used his quantity choices to signal his rival that they ought to coordinate, and remarkably, several of the pairs, with no more communication than this, were successful in alternating their play in this way. Appendix A provides data from this experiment, as well as a discussion of the various patterns of signalling the subjects used, and their degree of success.

Finally, the three learning models' predictions of subjects' period-by-period play were not supported by the experimental data: two full-distribution tests soundly rejected all three models as the process by which the data were generated. Thus, while the equilibria that are stable according to each of these three dynamic models (and stable under many other dynamics as well) are the ones that were actually played, our experiments provide no evidence suggesting which dynamic models are reasonable descriptions of empirical out-of-equilibrium behavior. This too is clearly an important subject for further research.

Acknowledgements

Financial support from the National Science Foundation (SBR-9311455) and research support services provided by the University of Arizona's Economic Science Laboratory (ESL) are gratefully acknowledged. Jason Shachat provided invaluable research assistance and many helpful comments. Zhu Li and Richard Kaiser, of the ESL, provided programming assistance.

Appendix A

Experiment UDF

As described in Section 7, in the fixed-matching experiment UDF the most profitable behavior for a pair of subjects was to coordinate their play, with only one subject in each period choosing a positive quantity. Some of the subject pairs successfully accomplished such coordination, and in several other pairs one of the subjects seems to have attempted to choose quantities that would lead his or her rival to cooperate, but without success.

Several patterns of successful and unsuccessful quantity signalling can be identified in the detailed data that appear below.

The pair of subjects #1 and #11 coordinated as follows. Following 14 periods of moderate gains and losses, subject #11 entered a zero quantity in period 15; this yielded a profit of 162 for subject #1 who had chosen 9 units. Subject #11 then sold 14 units in period 16 but subject #1 again sold 9 units, which yielded losses for both subjects. Subject #11 then continued to alternate between selling 0 units and 14 units, with no reciprocal choices by subject #1 until period 20. From period 21 to the last period, 37, subjects #1 and #11 successfully coordinated. They, and the other subjects, had not been informed of the number of periods in the experiment.

Subjects #2 and #12 had a relatively easy time coordinating. After losing money in 8 out of the first 9 periods, subject #12 chose zero units in period 10. This yielded a profit of 148 for subject #2, who sold 2 units. Subject #2 immediately reciprocated by choosing zero units in period 11 and the coordination was established.

Subjects #3 and #13 coordinated as follows. Following 14 periods of moderate gains and losses, subject #13 began in period 15 to alternate between 15 and 0 units. Subject #3 did not reciprocate in periods 15–17; then entered zero units in period 18 when subject #13 also chose zero. Quantities of (9, 15) yielded losses for both subjects in period 19. Then, beginning in period 20, subjects #3 and #13 coordinated their choices.

Subject #15 began to alternate between 15 and 0 units in period 21. By period 27, subject #5 had not reciprocated. Following quantity pairs of (15, 15) for the two subjects in periods 27–29, subject #15 sold 1500 units in period 30. This produced no response from subject #5 in periods 31 and 32. Then in period 33, subject #15 sold 18,928 units. This message appeared to get through to subject #5 and coordination was established in period 34.

The choices of subjects #4 and #14 show what appears to be a failed attempt at coordination. Subject #4 alternated between 0 and 15 units in periods 12–14. This produced no response from subject #14 who routinely sold 13 units. Subject #4 then sent a strong message by selling 16,959 units in period 15. Both subjects then chose 0 units in period 16. The quantity pair of (15, 0) might have started a coordination. But when subject #14 switched from 0 to 15 in period 18, subject #4 repeated her quantity choice of 15 rather than switching to 0. Subject #4 then repeated the strong message of selling 16,959 units in period 19, but followed this with a choice of 15 units in period 20 rather than the 0 units chosen in period 16 after the preceding strong message. Choice of 15 units by both subjects in period 20 produced losses. Subject #4 then repeated the 16,959 message in period 21. Subjects #4 and #14 never successfully coordinated, although it appears they may have achieved coordination just as the experiment was ending. When subject #4 was called up individually to be paid (a small amount) at the end of the experiment, she was very angry with her paired subject (whose identity was unknown to her).

Subjects #10 and #20 converged to the boundary Nash equilibrium of (0, 15) in period 14. Subject pairs #6 and #16, #7 and #17, #8 and #18, and #9 and #19 did not seem to find any pure strategy equilibria. It appears that subject pair #8 and #18 may have been about to coordinate when the experiment ended (Table 6).

Table 6
Experiment UDF: Fixed Pairs (Action, Opponent's Action, and Payoff for Each Subject in Each Period)

t	Subject #1			Subject #2			Subject #3			Subject #4			Subject #5		
	Act	Opp	Pay	Act	Opp	Pay									
1	5	12	-20	10	2	135	15	8	-24	4	7	26	15	8	-24
2	2	15	-17	8	10	-12	10	7	35	9	8	18	8	10	-12
3	9	14	-30	10	14	-30	11	9	-11	4	4	50	2	5	23
4	1	4	14	10	10	-25	7	8	21	14	4	102	3	10	3
5	4	12	-14	4	10	2	9	15	-30	5	4	60	10	10	-25
6	12	12	-30	10	10	-25	10	7	35	13	4	104	18	11	-30
7	1	6	10	10	7	35	9	9	0	9	13	-30	7	10	-7
8	1	5	12	10	10	-25	7	9	7	0	13	0	5	15	-30
9	2	14	-13	10	15	-30	8	9	4	10	13	-30	5	0	100
10	1	12	-2	8	0	148	5	8	20	10	13	-30	10	10	-25
11	12	12	-30	0	15	0	6	8	21	0	13	0	11	10	-30
12	11	5	77	15	0	216	7	7	35	0	13	0	2	10	3
13	10	5	75	0	15	0	8	7	36	15	13	-30	5	0	100
14	10	5	75	15	0	216	9	7	36	0	13	0	13	14	-30
15	9	0	162	0	15	0	9	15	-30	16959	13	-30	15	14	-30
16	9	14	-30	15	0	216	9	0	162	0	0	0	5	14	-30
17	8	0	148	0	15	0	9	15	-30	15	0	216	15	0	216
18	9	14	-30	15	0	216	0	0	0	15	15	-30	18	15	-30
19	15	0	216	0	15	0	9	15	-30	16959	15	-30	18	15	-30
20	0	14	0	15	0	216	15	0	216	15	15	-30	17	15	-30
21	15	0	216	0	15	0	0	15	0	16959	15	-30	19	15	-30
22	0	15	0	15	0	216	15	0	216	15	13	-30	19	0	184
23	15	0	216	0	15	0	0	15	0	0	0	0	20	15	-30
24	0	15	0	15	0	216	15	0	216	15	0	216	18	0	198
25	15	0	216	0	15	0	0	15	0	15	15	-30	15	15	-30
26	0	15	0	15	0	216	15	0	216	15	15	-30	15	0	216
27	15	0	216	0	15	0	0	15	0	15	15	-30	15	15	-30
28	0	15	0	15	0	216	15	0	216	15	15	-30	15	15	-30
29	15	0	216	0	15	0	0	15	0	15	15	-30	15	15	-30
30	0	15	0	15	0	216	15	0	216	15	15	-30	15	1500	-30
31	15	0	216	0	15	0	0	15	0	15	15	-30	15	15	-30
32	0	15	0	15	0	216	15	0	216	15	15	-30	15	15	-30
33	15	0	216	0	15	0	0	15	0	15	15	-30	15	18928	-30
34	0	15	0	15	0	216	15	0	216	0	15	0	0	15	0
35	15	0	216	0	15	0	0	15	0	15	15	-30	15	0	216
36	10	15	0	15	0	216	15	0	216	15	15	-30	0	15	0
37	15	0	216	0	15	0	0	15	0	15	0	216	15	0	216

Table 6 (Continued)

t	Subject #6			Subject #7			Subject #8			Subject #9			Subject #10		
	Act	Opp	Pay	Act	Opp	Pay									
1	16	13	-30	5	2	80	5	7	30	7	10	-7	10	7	35
2	11	13	-30	7	10	-7	7	8	21	10	16	-30	16	14	-30
3	5	13	-30	5	7	30	9	5	72	9	5	72	8	0	148
4	0	9	0	6	4	69	6	10	-3	9	5	72	9	1	144
5	12	9	-18	6	9	9	6	7	33	9	5	72	7	3	91
6	5	9	10	10	8	15	10	9	-5	9	1	144	11	3	121
7	11	9	-11	9	8	18	10	7	35	9	1	144	13	0	208
8	4	9	10	9	7	36	9	4	90	9	10	-18	13	15	-30
9	4	9	10	10	1	155	6	4	69	8	10	-12	15	14	-30
10	5	9	10	11	5	77	7	6	49	10	9	-5	11	13	-30
11	5	9	10	11	0	187	5	5	50	8	9	4	0	14	0
12	5	9	10	12	7	30	5	3	70	9	9	0	0	10	0
13	5	9	10	12	0	198	8	9	4	8	3	100	1	15	-8
14	5	9	10	12	0	198	9	5	72	8	7	36	0	15	0
15	4	9	10	13	5	78	8	5	68	7	4	77	0	15	0
16	4	9	10	14	0	214	10	8	15	7	6	49	0	15	0
17	6	9	9	14	0	214	10	10	-25	6	6	45	0	15	0
18	5	9	10	15	0	216	6	6	45	5	6	40	0	15	0
19	4	9	10	15	15	-30	7	6	49	6	6	45	0	15	0
20	5	9	10	14	15	-30	8	5	68	6	6	45	0	15	0
21	4	9	10	15	15	-30	8	8	20	6	9	9	0	15	0
22	5	9	10	15	15	-30	6	8	21	10	6	55	0	15	0
23	4	9	10	15	15	-30	9	8	18	6	6	45	0	15	0
24	5	9	10	13	15	-30	9	8	18	6	5	57	0	15	0
25	4	11	-6	10	15	-30	10	9	-5	5	5	50	0	15	0
26	2	9	7	8	50	-30	10	9	-5	5	5	50	0	15	0
27	4	11	-6	7	4	77	9	10	-18	5	5	50	0	15	0
28	2	9	7	7	7	35	5	10	0	5	4	60	0	15	0
29	1	11	0	8	5	68	10	10	-25	5	4	60	0	15	0
30	4	12	-14	8	5	68	10	10	-25	5	5	50	0	15	0
31	9	11	-30	9	5	72	10	15	-30	5	5	50	0	15	0
32	4	9	10	9	5	72	10	15	-30	5	5	50	0	15	0
33	1	9	4	9	9	0	15	10	-30	5	5	50	0	15	0
34	4	11	-6	8	15	-30	10	15	-30	5	5	50	0	15	0
35	9	11	-30	7	5	63	10	0	175	5	9	10	0	15	0
36	1	11	0	8	6	52	6	15	-30	10	5	75	0	15	0
37	4	11	-6	8	15	-30	15	0	216	5	5	50	0	15	0

Table 6 (Continued)

t	Subject #11			Subject #12			Subject #13			Subject #14			Subject #15		
	Act	Opp	Pay	Act	Opp	Pay	Act	Opp	Pay	Act	Opp	Pay	Act	Opp	Pay
1	12	5	42	2	10	-3	8	15	-30	7	4	56	8	15	-30
2	15	2	117	10	8	-15	7	10	-28	8	9	-20	10	8	-15
3	14	9	-30	14	10	-30	9	11	-30	4	4	38	5	2	65
4	4	1	62	10	10	-30	8	7	12	4	14	-30	10	3	85
5	12	4	66	10	4	65	15	9	-30	4	5	30	10	10	-30
6	12	12	-30	10	10	-30	7	10	-28	4	13	-30	11	18	-30
7	6	1	87	7	10	-28	9	9	-27	13	9	-30	10	7	5
8	5	1	75	10	10	-30	9	7	9	13	0	169	15	5	27
9	14	2	119	15	10	-30	9	8	-9	13	10	-30	0	5	0
10	12	1	138	0	8	0	8	5	44	13	10	-30	10	10	-30
11	12	12	-30	15	0	177	8	6	28	13	0	169	10	11	-30
12	5	11	-25	0	15	0	7	7	14	13	0	169	10	2	105
13	5	10	-15	15	0	177	7	8	0	13	15	-30	0	5	0
14	5	10	-15	0	15	0	7	9	-14	13	0	169	14	13	-30
15	0	9	0	15	0	177	15	9	-30	13	16959	-30	14	15	-30
16	14	9	-30	0	15	0	0	9	0	0	0	0	14	5	35
17	0	8	0	15	0	177	15	9	-30	0	15	0	0	15	0
18	14	9	-30	0	15	0	0	0	0	15	15	-30	15	18	-30
19	0	15	0	15	0	177	15	9	-30	15	16959	-30	15	18	-30
20	14	0	175	0	15	0	0	15	0	15	15	-30	15	17	-30
21	0	15	0	15	0	177	15	0	177	15	16959	-30	15	19	-30
22	15	0	177	0	15	0	0	15	0	13	15	-30	0	19	0
23	0	15	0	15	0	177	15	0	177	0	0	0	15	20	-30
24	15	0	177	0	15	0	0	15	0	0	15	0	0	18	0
25	0	15	0	15	0	177	15	0	177	15	15	-30	15	15	-30
26	15	0	177	0	15	0	0	15	0	15	15	-30	0	15	0
27	0	15	0	15	0	177	15	0	177	15	15	-30	15	15	-30
28	15	0	177	0	15	0	0	15	0	15	15	-30	15	15	-30
29	0	15	0	15	0	177	15	0	177	15	15	-30	15	15	-30
30	15	0	177	0	15	0	0	15	0	15	15	-30	1500	15	-30
31	0	15	0	15	0	177	15	0	177	15	15	-30	15	15	-30
32	15	0	177	0	15	0	0	15	0	15	15	-30	15	15	-30
33	0	15	0	15	0	177	15	0	177	15	15	-30	18928	15	-30
34	15	0	177	0	15	0	0	15	0	15	0	177	15	0	177
35	0	15	0	15	0	177	15	0	177	15	15	-30	0	15	0
36	15	0	177	0	15	0	0	15	0	15	15	-30	15	0	177
37	0	15	0	15	0	177	15	0	177	0	15	0	0	15	0

Table 6 (Continued)

t	Subject #16			Subject #17			Subject #18			Subject #19			Subject #20		
	Act	Opp	Pay												
1	13	16	-30	2	5	17	7	5	42	10	7	5	7	10	-28
2	13	11	-30	10	7	5	8	7	12	16	10	-30	14	16	-30
3	13	5	39	7	5	42	5	9	-5	5	9	-5	0	8	0
4	9	0	135	4	6	22	10	6	25	5	9	-5	1	9	1
5	9	12	-30	9	6	27	7	6	28	5	9	-5	3	7	12
6	9	5	45	8	10	-30	9	10	-30	1	9	1	3	11	-12
7	9	11	-30	8	9	-20	7	10	-28	1	9	1	0	13	0
8	9	4	63	7	9	-14	4	9	-2	10	9	-30	15	13	-30
9	9	4	63	1	10	-1	4	6	22	10	8	-15	14	15	-30
10	9	5	45	5	11	-25	6	7	15	9	10	-30	13	11	-30
11	9	5	45	0	11	0	5	5	35	9	8	-9	14	0	175
12	9	5	45	7	12	-30	3	5	24	9	9	-27	10	0	145
13	9	5	45	0	12	0	9	8	-9	3	8	6	15	1	147
14	9	5	45	0	12	0	5	9	-5	7	8	0	15	0	177
15	9	4	63	5	13	-30	5	8	5	4	7	14	15	0	177
16	9	4	63	0	14	0	8	10	-30	6	7	15	15	0	177
17	9	6	27	0	14	0	10	10	-30	6	6	27	15	0	177
18	9	5	45	0	15	0	6	6	27	6	5	39	15	0	177
19	9	4	63	15	15	-30	6	7	15	6	6	27	15	0	177
20	9	5	45	15	14	-30	5	8	5	6	6	27	15	0	177
21	9	4	63	15	15	-30	8	8	-4	9	6	27	15	0	177
22	9	5	45	15	15	-30	8	6	28	6	10	-21	15	0	177
23	9	4	63	15	15	-30	8	9	-20	6	6	27	15	0	177
24	9	5	45	15	13	-30	8	9	-20	5	6	25	15	0	177
25	11	4	66	15	10	-30	9	10	-30	5	5	35	15	0	177
26	9	2	99	500	8	-30	9	10	-30	5	5	35	15	0	177
27	11	4	66	4	7	14	10	9	-30	5	5	35	15	0	177
28	9	2	99	7	7	14	10	5	45	4	5	30	15	0	177
29	11	1	132	5	8	5	10	10	-30	4	5	30	15	0	177
30	12	4	66	5	8	5	10	10	-30	5	5	35	15	0	177
31	11	9	-30	5	9	-5	15	10	-30	5	5	35	15	0	177
32	9	4	63	5	9	-5	15	10	-30	5	5	35	15	0	177
33	9	1	117	9	9	-27	10	15	-30	5	5	35	15	0	177
34	11	4	66	15	8	-30	15	10	-30	5	5	35	15	0	177
35	11	9	-30	5	7	15	0	10	0	9	5	45	15	0	177
36	11	1	132	6	8	3	15	6	-3	5	10	-15	15	0	177
37	11	4	66	15	8	-30	0	15	0	5	5	35	15	0	177

References

- Brown, W.J., 1951. Iterative solution of games by fictitious play. In: Koopmans, T., (Ed.), *Activity Analysis of Production and Allocation*, Wiley.
- Conover, W.J., 1980. *Practical Nonparametric Statistics*, Wiley, New York.
- Epps, T.W., Singleton, Kenneth J., 1986. An omnibus test for the two-sample problem using the empirical characteristic function. *Journal of Statistical Computation and Simulation* 26, 177–203.
- Holt, Charles A., 1995. Industrial organization: A survey of the results of laboratory experiments. In: Roth A., Kagel, J., (Ed.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.
- Milgrom, Paul, Roberts, John, Adaptive and sophisticated learning in normal form games, *Games and Economic Behavior*, February 1991.
- Moreno, Diego, Walker, Mark, 1991. On recursive learning in noncooperative games. Discussion paper #91-20, Department of Economics, University of Arizona, Tuscon, AZ.
- Moreno, Diego, Walker, Mark, Two problems in applying Ljung's 'projection algorithms' to the analysis of decentralized learning, *Journal of Economic Theory*, April 1994.